BYTE

MAY 1990 A McGRAW-HILL PUBLICATION

PRODUCT FOCUS: THE TOP SQL SERVERS

Mainstream Amiga

New Amiga 3000: Mac II power, multitasking, multimedia... at a Commodore price

Desktop Supercomputing

Техника **PAGE 262**

The First 33-MHz 486s

New Battery Technology for Laptops

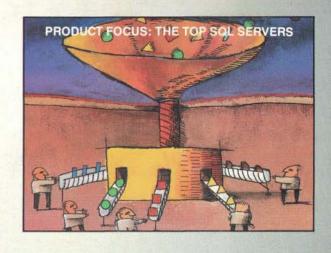**Smoother VGA**

Inside Tape Backup Systems

REVIEWS
4 386SX Laptops
5 New Trackballs
7 Updated TIGA Boards
Opus Desktop Tower

THE CREDENTIAL
A WORLD LEAD

Before the Dell System 325 was named the top 386 personal computer by the rest of the world, it had quite a following at home.

Winning both the *PC Magazine* Editor's Choice and *PC World's* Best Buy awards.

And most recently, it was rated number one for overall customer satisfaction, in a *PC Week* poll of corporate volume buyers.

One reason for this unprecedented popularity is that Dell deals directly with every Dell customer in the world.

Custom configuring each computer. Responding to the needs of each user.

So when you call Dell for a System 325, you get more than the exhilaration of running the world's top rated 386 computer, you get something called satisfaction.
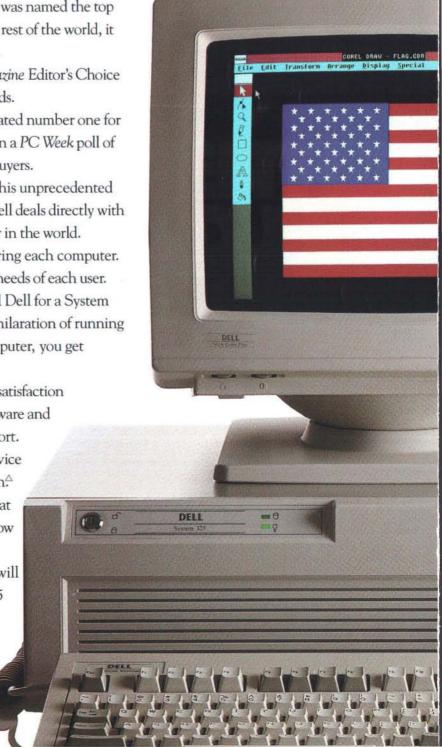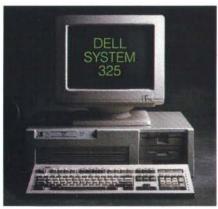
With a 30-day money-back satisfaction guarantee, self-diagnostic software and toll-free expert technical support. As well as next day on-site service provided by Xerox Corporation.△

And you can get it starting at $3,999. With leasing plans as low as $145 a month.**

Call us toll-free. And Dell will custom configure a System 325 for you.

After all, that's what the world is coming to.

THESE COUNTRIES JUST GOT TOGETHER AND ELECTED A NEW LEADER.

When computer publications from nine countries got together and voted for the best 386™ computer, it marked a major shift in power.

They chose the Dell System® 325 head and shoulders over every other system in the world.

Turn the page and find out why.

800-365-1240

DELL
COMPUTER
CORPORATION

FOR NETWORK OR UNIX INFORMATION, CALL 800-678-UNIX.

Germany

Hungary

Yugoslavia

Poland

Japan

United Kingdom

Spain

Italy

Years From Now,
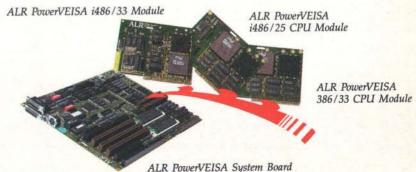
Standard Features at $5995...
EISA
33-MHz 386
64-KB Cache
5-MB RAM
106-MB Hard Disk
What more could you want?

An Honest, Affordable Upgrade Path.

You know what you want from a 386™ PC. ALR's **PowerVEISA 33/386** delivers. Cached 33-MHz performance, plenty of RAM to run even the bulkiest of today's applications, a choice of high performance hard disks, and **EISA** compatibility — all at a price that puts many similarly equipped ISA systems to shame. So what more could you want?

How about the future?

ALR's PowerVEISA 33/386 offers you an honest, affordable upgrade path to

25-MHz, 33-MHz (available second quarter) and future i486™ processors.

ALR PowerVEISA i486/33 Module

ALR PowerVEISA i486/25 CPU Module

ALR PowerVEISA 386/33 CPU Module

ALR PowerVEISA System Board

CONTENTS

May 1990
Volume 15, Number 5

Desktop Supercomputing/204

COVER STORY

**First Impressions:
Commodore Sets Course
for Multimedia
page 122**

Commodore's
Amiga 3000 has
a lot to offer
for a modest price.

COVER PHOTOGRAPHY: PAUL AVIS © 1990

Around the World in Text Displays/262

Expert Advice—Macinations/103

DESKTOP SUPERCOMPUTERS AND Техника

Two researchers discover supercomputer power that's free for the taking, and we discover a staff member's hidden past

Dave Gelernter and Jim Philbin have a radical mission: They want to make parallel programming as easy as conventional programming. They believe that with Linda—a "coordinating language" that can provide parallel-processing capabilities to any conventional language— they can accomplish their goal.

Dave Gelernter conceived of Linda while he was a graduate student at the State University of New York at Stonybrook in 1982. In fact, Linda formed the basis of his Ph.D. thesis. The reaction to Linda from the parallel-processing community was immediate and unanimous. Says Gelernter, "People thought it was elegant and beautiful, and that it was totally unimplementable."

By the mid-1980s, that unanimity of opinion was shattered. "In late 1985, the first paper describing an actual Linda implementation appeared," says Gelernter, "and people were shocked."

After that, says Gelernter, Linda polarized the parallel-processing community. As he puts it, Linda advocates are arrayed against a host of "navel-gazing computer scientists out of touch with reality." Behind the humor lies the core premise of Linda. As Jim Philbin puts it, "You shouldn't need a Ph.D. to write a parallel program."

The reason behind the controversy is that Linda consists of just six statements, and yet can add parallel processing to any conventional language. "Most people in academia think you have to create parallel languages from scratch," says Gelernter. "Linda lets you use the language you're used to."

Gelernter and Philbin are betting that Linda will make a big splash in the commercial world. In the near term, that means Hypercomputer Linda and new C-Linda implementations for machines such as the Intel and NCube hypercubes. It also means new versions, such as Ada-Linda and C++ Linda.
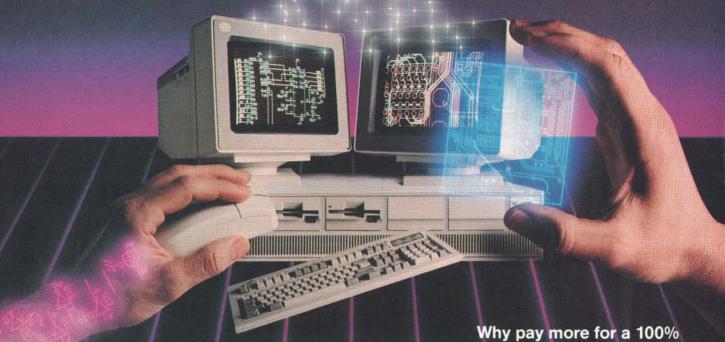
Both Gelernter and Philbin foresee a bright future, for Linda and for parallel processing in general. Gelernter sees a future where not only parallel processing, but also distributed databases and distributed processing using heterogeneous machines, will become much more widespread and important. Linda can handle all these things. As Gelernter puts it, "It's a boom time for the Linda business."

For more on Linda, see the article "Spending Your Free Time" by David Gelernter and James F. Philbin. It begins on page 213 of this issue.

Every once in a while, we discover something interesting that we hadn't known about one of our staff members. For example, our features editor recently decided to do an article about the problems involved with computing in non-roman languages (those that have different alphabets and writing styles from English). On checking with the staff, he found that technical editor Ben Smith had once been involved in a lengthy project to create an Arabic computer system. Ben is also the only person we know who "just happens" to have Cyrillic fonts on his home Macintosh. Indeed, Ben is *very* familiar with the issues involved in non-roman computing—probably more familiar than he ever wanted to be.

Ben's article in this issue, "Around the World in Text Displays," explores the challenges of computing in Cyrillic, Hebrew, Arabic, Chinese, and Japanese, and how the tasks are handled differently on a PC compatible and on a Macintosh. ∎

At last, an assistant that follows your directions

Wouldn't it be great to delegate your routing?

You can! We know your time is valuable. That's why Wintek pioneered comprehensive and affordable CAD packages for IBM personal computers. HiWIRE-Plus continued that tradition, integrating schematic-capture features and printed-circuit-artwork capabilities into one versatile package.

New autorouter.

The Autorouter for HiWIRE-Plus is powerful enough to handle the most demanding design problems, yet simple enough for a casual user. Just turn it loose on your design. It's hassle free because it works long hours, without supervision or errors.

100% autorouting.

The autorouter for HiWIRE-Plus rips-up, reroutes, and with appropriate design rules, racks up 100% completion.

☐ Forget gridded routers. This autorouter places vias and traces anywhere your design rules allow. With 1-mil resolution.

☐ Vary trace width and spacing for individual networks. Route 1, 2, 3, or more tracks between IC and connector pins.

☐ Set up boards from 1 to 250 layers, up to 60″ × 60″.

☐ Specify shape, size, and type of vias, layer-by-layer: through-hole, blind, buried, micro. Specify via types for individual networks.

☐ Use fewer vias and layers than comparably priced autorouters.

☐ For use on your IBM PC, XT, AT, PS/2, or compatible with 640K RAM.

Why pay more for a 100% autorouter?

Compare the features and performance to packages costing five times more. HiWIRE-Plus and the Autorouter for HiWIRE-Plus sell for $895 each. Both have a no-nonsense, 30-day money-back guarantee. With unlimited, toll-free, no-charge technical support.

Let HiWIRE convince you that it makes a great assistant. Call us toll-free at (800) 742-6809 today and put HiWIRE-Plus and the Autorouter for HiWIRE-Plus to work for you tomorrow.

WINTEK®

Wintek Corporation
1801 South Street
Lafayette, IN 47904-2993
Fax: (317) 448-4823
Phone: (317) 742-8428 or
(800) 742-6809

Until now, this is h
and word processo

Use Microsoft® Excel and Microsoft Word for Windows™ and you'll no longer be saddled with the problem of combining words, numbers and graphics. You see, they were designed to work together. Which means they share a good deal more than a name.

Like data for instance. Now you can cut and paste words, tables, charts, whatever. From any document, in either application, any which way.

PAPERLESS FAX, AND AN INVITATION FOR YOU

A flood of faxes and BYTE's upcoming 15th anniversary prompt some questions

I like fax machines. I use them a great deal, especially when I need to send something in a hurry to someone without a computer or who doesn't telecommunicate; or when I need to see or send hand drawings, and so forth. Faxes can be an invaluable way to move data around.

But in many instances, faxing makes no sense at all. People seem to fax for the same reason that others climb Mount Everest: just because it's there.

Fax Abuse

Here's an example of fax abuse: Person X uses a computer to create a text message, prints a hard copy, and feeds the paper into a fax machine. At the other end of the line, Person Y's fax churns out more paper. Then, in order for Person Y to do anything useful with the faxed information, he or she (or an unlucky secretary) must expensively rekey the data back into a computer.

Because the information starts and ends inside a computer, the two paper copies and the manual rekeying are wasteful, redundant, and just plain unnecessary.

Plugging a fax board into your computer can help avoid some of the paper-based steps, but this carries other problems. For example, faxes are bit-mapped images, so they're real space hogs: A two-page fax memo can easily consume more than 200,000 bytes of bit-map disk space. And, because it is a bit map, you still have to read the fax image off your screen and separately rekey the pertinent information by hand. What a waste.

Yes, there have been some promising advances involving fax boards and optical-character-recognition software, but we're not there yet. When bit map–to-ASCII software is more common and reliable, that two-page memo will occupy just 4K bytes of disk space, instead of 200K bytes.

A Dream Alternative

Now, imagine a reliable, paperless fax, a simple way to send normal documents in a compact form that can be piped directly into or out of your computer without wasting time or paper or disk space. The format would be standardized so that dissimilar computers (different hardware, different operating systems) still can work with it without hassles, letting anyone directly reuse data sent from anyone else in one step without any rekeying whatsoever.

This paperless faxing would be as convenient as regular faxing, and—let's pull out all the stops—it would also be dirt cheap. The hardware would start at well under $100, and paperless fax software would range from a few hundred dollars down to absolutely free. Paperless faxing obviously won't use any plain or expensive fax paper at all, and let's say that its telecommunications charges would be less than for standard faxing. Paperless fax users will win on every score.

OK, OK; you've figured it out. I'm talking about E-mail. Somehow, in all the hullabaloo about faxing, E-mail seems to have disappeared from some people's consciousness, and it's a shame.

Half the Cost

Take that standard 4000-character, two-page letter: If you fax it coast to coast during normal business hours, it will cost you about 50 cents to transmit—more or less, depending on the vagaries of your specific situation. If it's a paper memo sent to a paper-based fax machine, you'll also generate four pages of hard copy in the process, costing a few more cents. Or, if it's a bit-mapped fax, you

will consume over 200,000 bytes of disk space.

Now E-mail the exact same text: It will cost less than half as much to send, about 20 cents. You'll generate no paper; if captured to disk, you'll use just 4K bytes of disk space. Best of all, the information will be in a format you can do something with, instantly, on your computer, with no expensive (labor-intensive) rekeying. That's right: It costs less than half as much, and yet it's even more useful.

Faxing has its place—it's become a vital link in business communication. But when you're sending ordinary text printouts, do yourself, your pocketbook, and your recipients a favor: Try E-mail.

An Invitation

This September will mark the record-setting 15th year of BYTE's publication, and the BYTE editors are busy assembling an extraordinary anniversary issue for you.

We'd like you to help us write this special issue. Share your views with a world-wide audience of over a million and a half sophisticated computer users! Drop us a note and tell us what you think the best and worst developments of the last 15 years have been. What do you want your computer to be able to do 5, 10, or 15 years from now? Where do you think all this technology should lead us over the next few years? How can the computer industry better meet your needs and the needs of your company? What new technologies would you like to see developed? Any computer-related subject is fair game; we want to hear from you!

We plan to edit the replies we receive and print the best in the September issue. Please send your thoughts to me at the address shown on the masthead or via MCI Mail or BIX mail. And what the heck—if you must, you can even fax it.

—Fred Langa
Editor in Chief
(BIX name "flanga")

Better is

The new 16″ MultiSync® 4D and 20″ 5D both offer compatibility with

a wide range of graphics boards and computer systems.

NEC's award-winning multiple frequency technology, now available in larger, digitally controlled, high-resolution color monitors.
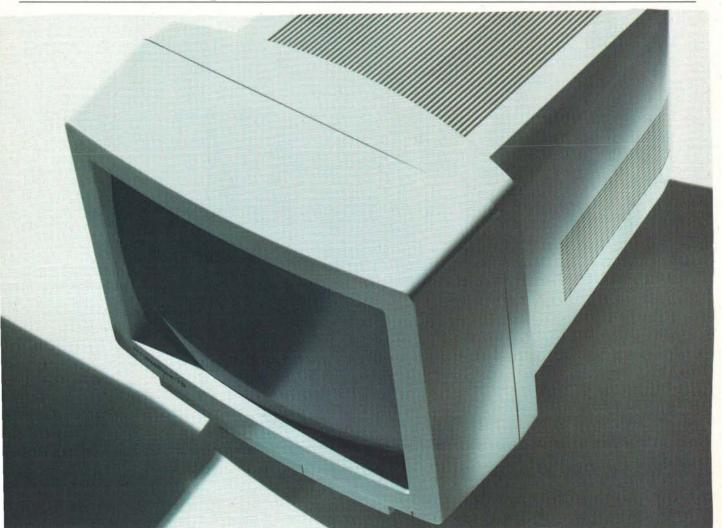
Both are optimized for the IBM PS/2, PC/AT/XT (and 100% compat-

ibles) and the Macintosh II. Both offer a microprocessor-based digital

control system for preset and custom graphics modes, automatic screen

C&C
Computers and Communications

bigger.

configuration and optimal image. And both give you great resolution—

the 4D, from VGA up to 1024 x 768; the 5D, from VGA to 1280 x 1024.

The main difference: one is for large ideas. The other, for extra-

large. For literature, call 1-800-826-2255. For details, call NEC

Technologies, Inc. 1-800-FONE-NEC.

NEC

Just when Power Seekers thought they had nowhere left to go...

Introducing the

IBM RISC System/6000™ family.

With the most powerful desktop workstation on earth.

No matter how big the job you're setting off to conquer, the new IBM RISC System/6000 family of POWERstations and POWERservers is for you. The desktop POWERstation 320 puts more than 7 MFLOPS and over 27 MIPS at your command—power that soars as high as 13 MFLOPS and 41 MIPS in other models.

What makes it possible is POWER Architecture—Performance Optimization With Enhanced RISC—IBM's second generation of RISC technology. It delivers up to four instructions per cycle and has the most powerful CMOS microprocessor in the business.

	MFLOPS	MIPS	SPECmark™
POWERstation 320[†]	7.4	27.5	22.3
DECstation™ 3100[‡]	1.6	14.3	10.1
SPARCstation™ 1[‡]	1.4	12.5	8.4

[†]MFLOPS are the results of the double-precision, all FORTRAN Linpack test 100 x 100 array suite. The Dhrystone Version 1.1 test results are used to compute RISC System/ 6000 Integer MIPS values where 1,757 Dhrystones/second is 1 MIPS (VAX 11/780). SPECmark is a geometric mean of ten benchmark tests.
[‡]Performance data are based on published benchmark information.

Micro Channel™ makes bottlenecks ancient history.
All RISC System/6000 models feature a new implementation of the IBM Micro Channel bus with I/O throughput of up to 40 megabytes per second. And the Micro Channel Architecture can accommodate a doubling or even quadrupling of this data transfer capacity, making traditional, nonexpandable architectures seem primitive by comparison.

Solid support for all UNIX® applications.
These systems are all based on the industry-standard UNIX operating system. And they'll run hundreds of applications in such diverse fields as engineering design, fluid dynamics, molecular modeling, securities trading, technical publishing and geophysical modeling, plus a wide selection of commercial applications.

*The POWERstation 320 comes with 8MB of RAM, 120MB of disk, 19-inch grayscale display and graphics adapter, keyboard, mouse, Ethernet™ adapter, AIX operating system. user interface environment, NFS™ and hypertext documentation search and retrieval capability.

A brave new world of 3D graphics. All our POWERstations are built to deliver high-speed, high-resolution graphics. Our 3D processors offer an almost unimaginable palette of 16 million colors. And our new Supergraphics POWERstation 730 gives you 3D graphics capabilities that are really out of this world, performing nearly one million 3D vector transformations per second.

AIX™ brings different worlds together.
Every model can come preloaded with AIX Version 3, IBM's version of the UNIX operating system. It supports major industry standards, like POSIX™ and TCP/IP, plus the programming languages C, FORTRAN, COBOL and Pascal. And AIX allows the RISC System/6000 family to fit into your installed UNIX base, as well as your IBM SAA™ environments, giving you the best of both worlds.

Supercomputer speed at a workstation price: $12,995.
Our entry desktop POWERstation 320 starts at $12,995 complete.* And every member of the RISC System/6000 family comes with software service at no extra charge and a 12-month warranty.

The worldwide support of IBM.
Then, of course, there's IBM service when you need it, 24 hours a day, 365 days a year. To find out more, call your IBM marketing representative or Business Partner. For literature, call 1 800 IBM-2468, ext. 225. And find out why it's the beginning of a new world for Power Seekers.
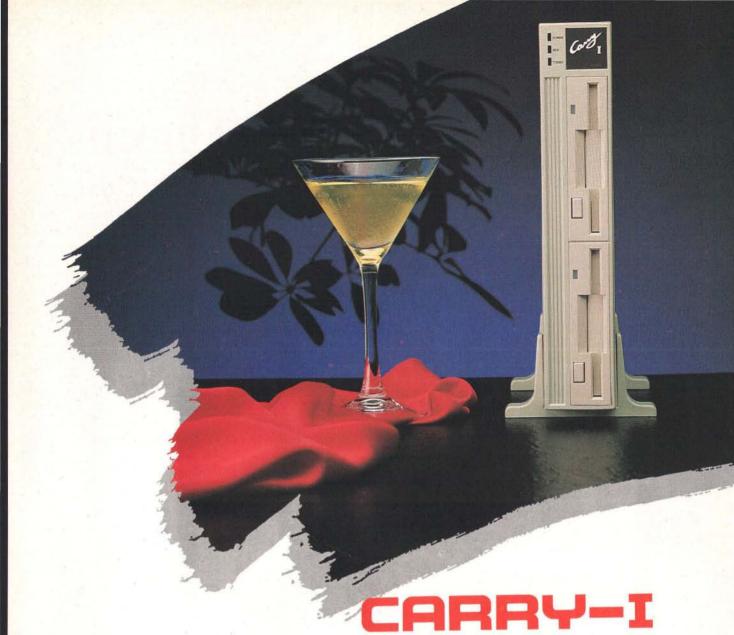
For the Power Seeker.

CARRY-I

The World's First & Original Book-Size Computer

FT-8100 Series
10MHZ XT/AMI BIOS/256K RAM expandable to 640k/One to two 720KB 3.5″ FDD/ Serial/Parallel/Game/CGA/MGA/Standard keyboard connector/16Watt Power adapter
Dimension: 240mm × 185mm × 45mm Weight: 1.9kg-2.4kg

FT-8200 Series
12MHZ, 0 Wait State AT/AMI BIOS with Diagnostic/1MB RAM/One to two 1.44MB 3.5″ FDD/Option 20MB or 40MB HDD/2 Serial/1 Parallel/CGA/MGA/Standard keyboard connector/30Watt Power adapter
Dimension: 240mm × 185mm × 45mm Weight: 2.1kg-2.8kg

FT-7082
82key mini-size Keyboard/XT-AT autoswitchable
Dimension: 310mm × 145mm × 27mm

It's priced lower than you'd expect for a PC with this kind of power and portability.
But don't take our word for it.
Call us today for more information.

COMDEX/Spring
Georgia World Congress Center
West Hall, Booth No: 4038

FLYTECH TECHNOLOGY CO., LTD.
(HEAD OFFICE)
2 FL., NO. 8, LANE 50, SEC. 3,
NAN-KANG RD., TAIPEI, TAIWAN, R.O.C.
TEL: (02)785-2556 FAX: (02) 735-2371, 783-7970
TELEX: 22233 FLTCO

FLYTECH TECHNOLOGY HANDELS-GMBH
MENDELSSOHNSTRASSE 53,
6000 FRANKFURT AM MAIN 1, WEST GERMANY
TEL: (069)746-081, 746-453 FAX: (069)749-375

FLYTECH TECHNOLOGY (U.S.A), INC.
3008 SCOTT BLVD., SANTA CLARA, CA. 95054 U.S.A.
TEL: (408)727-7373, 727-7374 FAX: (408)727-7375

FLYTECH TECHNOLOGY (H.K.) LTD.
B12, 8 FL., BLOCK B, TONIC INDUSTRIAL CENTRE,
19 LAMHING ST., KOWLOON BAY, KOWLOON,
HONG KONG
TEL: 305-1268 FAX: 796-8427

PC/XT/AT, are registered trademarks of International Business Machines Corp.   CARRY-I is a registered trademark of FLYTECH TECHNOLOGY CO., LTD.

MICROBYTES

Research news and industry developments shaping the world of desktop computing

Image Compression Chip Cuts Space and Time

In a development that could make graphics less of a hog when it comes to devouring computing resources, C-Cube Microsystems (San Jose, CA) has built a compression/decompression processor that can squeeze data, on the fly, by as much as 25 times. C-Cube, founded by former Weitek and Radius engineers, says that its custom VLSI processor can compress a 25-MB graphics image to 1 MB in 1 second or a 1-MB video frame in $\frac{1}{30}$ second, allowing for real-time compression of graphics images.

With this kind of compression power, the chip could significantly alleviate some of the processing and storage demands of such applications as multimedia, desktop video, and color desktop publishing. According to marketing manager Mauro Bonomi, there has been tremendous interest in the chip among manufacturers of computer equipment as well as makers of consumer electronics goods.

The single-chip CL550 uses the proposed Joint Photographic Experts Group (JPEG) standard compression algorithm, supported by IBM, DEC, and NEC. The basic compression technique involves discarding data in the graphics image describing frequencies that are invisible to the human eye. Thus, with compression ratios of 10 to 1 for screen images and 25 to 1 for print images, the difference in visual quality is hardly noticeable

(although quality degrades progressively as the compression ratio is increased).

The advent of data compression processors as standard equipment in computers and devices such as digital cameras and VCRs has major implications for graphics technology. One of the big stumbling blocks to transferring and manipulating high-resolution color images or full-motion video is the sheer size of the data files required to store those images. In 24-bit color, an 8½- by 11-inch image at 300 dpi requires 25 MB of storage space. One second of full-motion video at 30 frames per second requires as much as 30 MB of storage. The ability to squeeze the graphics data by as much as a factor of 25 means faster transmission rates and much smaller memory, storage, and bus bandwidth requirements.

The CL550 will ship to OEM customers later this year. The 27-MHz version costs $155 in quantities of 10,000, while the 10-MHz version is $95 in quantities of 10,000.

C-Cube is now selling a software version of the JPEG compression algorithm, called the Compression Workshop, for the Mac II for $550. A 30-to-1 compression of a 30-MB image takes about 1 hour using the software algorithm on the 68030 processor of the Mac IIx. It takes 1 second on the CL550, C-Cube says.
—*Nick Baran*

TrueType vs. Adobe: A Royal Mess for Users?

Now that Microsoft and Apple have released more details about their outline font technology for creating smooth, scalable characters on screen and paper (formerly called Royal and now called TrueType), the advantages are still not obvious and are months away from being real.

Many users are asking if there is something "better" about TrueType that makes it functionally better than Adobe fonts. Looking at the Adobe Type 1 font format and the technical

information available on TrueType, it's hard for users to see where either font design has a significant advantage over the other. Both use a set of points to describe the outline of a character. Both use "hints" (instructions that make characters look good, particularly at low resolution, such as on a monitor or impact printer), although they take different approaches. One apparent difference is that a portion of the TrueType hinting mechanism is
continued

Object-oriented programming was a popular topic at the recent Software Development '90 conference, judging by attendance at OOP seminars and the number of OOP products on display. (We saw demonstrations of several commercial applications created using SmallTalk, Actor, and Eiffel, but very few using C + +.) In talking to programmers, one theme emerged: Most of them want to learn about OOP but haven't made the further step of using OOP techniques professionally. There is considerable concern about compatibility between different object-oriented languages and environments.

Evidence at the Software Development conference is that **OS/2 is arriving**. Most of the major compiler manufacturers, with the notable exception of Borland, now have compilers that produce OS/2 code. **Watcom**, **JPI**, and **Lattice** all announced new compilers supporting OS/2. Conference sessions and classes covering OS/2 programming were also mostly sold out. Realia, which makes Cobol tools, is planning an OS/2 Presentation Manager version of its compiler.

Jensen and Partners International (Mountain View, CA), best known for its TopSpeed programming tools, plans **new Pascal and C + + compilers** for release this summer. TopSpeed Pascal will use the same environment as the company's Modula-2 system and will be fully ISO-compatible. It will have object-oriented extensions, including multiple inheritance. JPI will package it with a translator facility to make the compiler compatible with Borland's Turbo Pascal, as well as library extensions to provide Microsoft compatibility. Other libraries will provide a full DOS/BIOS interface, extended graphics, support for Borland's graphic interface (BGI), and a text window manager and time-sliced scheduler. Details of the C + + compiler were sketchy, but it will be in the same mold as other TopSpeed compilers.

independent of the character's outline, and it generally takes more points to describe a TrueType character. TrueType characters and Type 1 characters (rendered with Adobe Type Manager) look equally nice on the screen.

Microsoft and Apple say that their approach is faster at rendering characters, and some type designers who have been working with True-Type support that claim. (In a demo given by Microsoft, beta TrueType was noticeably faster than Adobe Type Manager at putting a PageMaker document on the screen.) Regardless, neither type of font will be rendered often. The computer or printer will construct a bit map from the font format at the proper resolution and then cache it in memory, using these bit maps until the memory is required for a new typeface or a graphic. A small difference in speed between the Adobe and TrueType interpreters probably won't amount to much, since they won't have to do their jobs frequently. This problem might be more acute on a computer system where the interpreter is displaying typefaces on a computer screen, but a large RAM font cache can eliminate frequent reconstructions of bit maps.

Microsoft and Apple promised originally that their new outline fonts would be part of the Mac's System 7.0 and OS/2 2.0, both due later this year. Microsoft announced recently that TrueType will also be part of its Windows environment sometime this year—not in the "imminently

coming" version (the legendary 3.0), but in what will probably be Windows 3.1, said Steve Ballmer, Microsoft vice president of system software. ("If not this year, then next," he said.)

Many of the good things that Microsoft and Apple claim about TrueType are of more benefit to type designers than to computer users who aren't typographers. They say that TrueType fonts are easier to develop; it's easier to convert fonts to the TrueType format than to Adobe Type 1 format; and because TrueType is similar to a low-level programming language, it's more flexible, and hence more appropriate for nonroman characters (kanji, in particular) than a high-level language like Adobe's. In a random sampling of people who work on PostScript fonts and rasterizers, we found many who agreed with the claims about TrueType. However, we also found many who said that they'll stick with the Adobe format because it's proven, it's established, and in some ways it's better.

But right now, all we can compare is something that exists—Adobe's PostScript fonts—against what Apple and Microsoft say will exist. What Apple and Microsoft have to do is deliver something that works—and works reliably. Microsoft's Ballmer agrees. "The most important thing is for us to show this stuff working across all the promised platforms," he said. "That's the whole point of personal computers: Things should just work."

—Tom Thompson and D. Barker

Apple Finally Gives Buyers a One-Year Warranty

I t's not quite the Berlin Wall tumbling down, but for buyers of Apple hardware, it's a triumphant moment. Responding to the many critics of its service policies, Apple Computer (Cupertino, CA) has finally changed its warranty. The company announced that all Apple hardware products sold in the U.S. after January 1, 1990, will be covered by a one-year limited warranty. Until now, Apple's often-bashed warranty has covered 90 days. The new limited warranty also provides for U.S. customers to have service performed anywhere in the world. Apple dealers outside the U.S. already offer a one-year warranty.

Customers who bought Apple

hardware in the U.S. prior to January 1, 1990, will get six months of AppleCare (Apple's program for out-of-warranty service) free if they contract (between March 19 and May 31) for six months of service.

The promotional price for a year of AppleCare for a Mac SE (with an internal 20-MB hard disk drive) is $162; for a Mac IIcx (with an internal 40-MB hard disk drive and a color monitor), the price is $246. Volume discounts of up to 25 percent are also available, depending on the number of Apple products covered.

Most retailers said that the new policy will help sell Apple equipment. One dealer said, "I've always had to

Until now there was only one way to integrate C and Assembler.

While C and Assembler give you power to burn, switching back and forth between them can leave your brain feeling a little fried.

All that stopping. And starting. And constantly retracing your steps.

Well, relax. Now there's Microsoft® QuickAssembler. Available with our clever QuickC® Compiler in one location: the first integrated environment for C and Assembler.

For the first time, you can save time with an integrated editor, compiler, assembler and debugger that let you create C programs, mixed C and Assembler programs, or Assembler programs that stand alone.

To make sure you feel at home in your new environment, we've designed Microsoft Quick Advisor, a hypertext electronic manual that coaches, coaxes and guides you on screen.

Quick Advisor gives you access to information on all ROM BIOS and MS-DOS® calls. And it even lets you cut and paste sample programs, so you can make both C and Assembler subroutines part of your routine in no time.

For more details on the incredible integrated power of QuickAssembler and QuickC Compiler, call (800) 426-9400. If you own QuickC Compiler version 2.0 already, we'll tell you how to add on QuickAssembler quick.

And take a load off your mind.

Microsoft®
Making it all make sense.™

PostScript® laser printers used to be expensive, large and shared by the entire office.

Until now.

Claris (Santa Clara, CA) has acquired from Red Brick Systems a **query technology** that's supposed to make it easier for Macintosh users to access data from large corporate databases. The Executive Query Tool lets Mac users log onto relational databases using the standard Mac interface and download data without having to learn the SQL programming language. When a user on a client-server network makes a query, EQT generates the SQL code, which then queries the database. Claris plans to develop a new application based on EQT by the end of this year. EQT runs on top of standard server-access protocols like Apple's CL/1, OCI from Oracle, and DB-Lib from Sybase, which will allow for access to most corporate databases, Claris claims.

IBM's new **RISC System/6000** appears to be attracting significant **support from Unix software vendors**. Over 70 software companies exhibited products at IBM's San Francisco rollout of the new machines. While most of the demonstrated applications were for scientists and engineers, there were also database applications from Oracle, Ingres, and Informix, as well as financial and business management software. The programs shown were running under a prerelease version of AIX 3. Most used the X Window System and some of the preliminary graphics objects supplied in OSF/Motif. According to a list from IBM, about 200 companies say they'll be bringing their programs to the RS/6000.

The new CD-ROM presentation of Mozart's *The Magic Flute*, from **Warner New Media** (Burbank, CA), couples more than 7000 screens of information with 143 minutes of digital music. A HyperCard stack supplies extensive annotation, commentary, musical examples, extra narration, and excerpts from other recordings, as well as a full index. The three-disk set will sell for a suggested $66. You need a 1-MB Mac and Apple CD-ROM drive.

very carefully explain Apple's warranty terms to customers, but not until they have already decided on Apple equipment. If I do it sooner, they ask why it isn't longer than 90 days. But I must say that the 90-day warranty has always been a great reason to then sell AppleCare."

The warranty was just one item of good news for users. Apple cut suggested retail prices of the Mac SE and SE/30 by as much as 17 percent and similarly shaved prices of the LaserWriter IINT and IINTX. And as of March 1, the price of Apple's CD SC CD-ROM player is $899, 25 percent less than it used to be.

—Larry Loeb

Tools Will Give Windows a Role in Multimedia

I f multimedia computing were a Hollywood movie, the Macintosh and the Amiga would be the stars, the De Niro and Streep. The IBM PC compatible would be the executive producer, running spreadsheets and taking care of unglamorous business chores.

But PC compatibles could grab a piece of the multimedia limelight this year, if developers come through with graphics, video, and audio tools for Microsoft Windows. Microsoft and other software vendors have now pledged to develop Windows-based tools for putting together multimedia presentations. These tools would be for recording, encoding, and editing audio; scanning, drawing, painting, and manipulating photos and graphics; capturing, editing, and digitizing motion video and animation; and editing and indexing text. Some of these packages are already working with the Mac, which suggests compatibility between Mac and Windows applications.

Microsoft has licensed Macro-Mind's multimedia playback software, Director, for use in a future version of Windows. Director 2.0 produces animation sequences for the Mac II in full color and with sound. The new Windows version (whose release date has not yet been announced) will be able to play back multimedia presentations prepared on the Macintosh. The same Macintosh files can be played through Windows, the companies said. (They recently demonstrated an IBM PC and Mac side by side playing Director 2.0 presentations.) Macro-Mind and Microsoft have agreed to explore additional cooperative efforts in related areas, including OS/2-based versions of animation run-time software.

"For multimedia to really take off, the industry needs a rich set of software tools from vendors across a wide range of disciplines," said Microsoft chairman Bill Gates.

Among the companies that are pledging to ship such tools for Windows-based applications by the end of this year are Asymetrix, Autodesk, Corel Systems, Electronic Arts, Farallon Computing, Zsoft, Meridian Data, and Micrografx.

—David Reed

Will This Be The Year of the Wireless LAN?

W e've been hearing for years about "The Year of the LAN." But at the recent ComNet show in Washington, D.C., people were talking instead about "The Year of the Wireless LAN." For users who don't need the high data rates offered by digital land lines or who want to avoid the hassles of cabling, several companies are, or soon will be, offering systems for wireless communications. Each uses spread-spectrum technology, which employs high-frequency radio signals to link computers. This approach allows signals in a narrow frequency band to be spread and transmitted over a broad range of frequencies with lower

energy content, thus minimizing noise and interference with other radio devices. The signal is then collapsed back to its original narrow frequency band at the receiving end of the transmission.

O'Neill Communications (Princeton, NJ) was one of the first to take this approach, with its LAWN (Local Area Wireless Network), a proprietary LAN that uses small transceivers to connect as many as 20 personal computers in a workgroup.

Ferranti Datacom (Sunnyvale, CA) plans to soon release the RM38, a $3450 radio modem designed for

Thousands Of People Do This Every Week...

12 MHZ 286 VGA

- 80286-12 Processor
- 2 Megs RAM
- 1.2 Meg 5¼" Drive
- 1.44 Meg 3.5" Drive
- 65 Meg 28ms RLL Drive
- 16 Bit VGA with 512K
- 14" 1024 x 768 Color Monitor
- 1 Parallel & 2 Serial Ports
- 101 Key Keyboard
- MS DOS 3.3 or 4.01

$1995.00

GATEWAY 386SX

- 2 Megs RAM
- 1.2 Meg 5¼" Drive
- 1.44 Meg 3.5" Drive
- 65 Meg 28ms RLL Drive
- 16 Bit VGA with 512K
- 14" 1024 x 768 Color Monitor
- 1 Parallel & 2 Serial Ports
- 101 Key Keyboard
- MS DOS 3.3 or 4.01

$2195.00

INFO WORLD — EXCELLENT VALUE — Jan. 15, 1990

20 MHZ 386 VGA

- 4 Megs RAM
- 1.2 Meg 5¼" Drive
- 1.44 Meg 3.5" Drive
- 65 Meg 28ms RLL Drive
- 16 Bit VGA with 512K
- 14" 1024 x 768 Color Monitor
- 1 Parallel & 2 Serial Ports
- 101 Key Keyboard
- MS DOS 3.3 or 4.01

$2595.00
With 64K Cache $2895.00

1988 BYTE AWARD OF DISTINCTION

25 MHZ 386 VGA

- 4 Megs RAM
- 1.2 Meg 5¼" Drive
- 1.44 Meg 3.5" Drive
- 110 Meg ESDI Drive
- ESDI Controller with 32K Cache
- 16 Bit VGA with 512K
- 14" 1024 x 768 Color Monitor
- 1 Parallel & 2 Serial Ports
- 101 Key Keyboard
- MS DOS 3.3 or 4.01

$3095.00
With 64K Cache $3495.00

33 MHZ 386 VGA

- 64K Cache RAM
- 4 Megs RAM
- 1.2 Meg 5¼" Drive
- 1.44 Meg 3.5" Drive
- 150 Meg ESDI Drive
- ESDI Controller with 32K Cache
- 16 Bit VGA with 512K
- 14" 1024 x 768 Color Monitor
- 1 Parallel & 2 Serial Ports
- 101 Key Keyboard
- MS DOS 3.3 or 4.01

$3995.00

PC WORLD — BEST BUY

25 MHZ 486 VGA

- 4 Megs RAM
- 1.2 Meg 5¼" Drive
- 1.44 Meg 3.5" Drive
- 150 Meg ESDI Drive
- ESDI Controller with 32K Cache
- 16 Bit VGA with 512K
- 14" 1024 x 768 Color Monitor
- 1 Parallel & 2 Serial Ports
- 101 Key Keyboard
- MS DOS 3.3 or 4.01

$5295.00

The 17-member Coordinating Committee for Multilateral Export Controls (COCOM) has agreed to **reduce restrictions on high-technology exports to the Eastern Bloc** countries undergoing political change—Czechoslovakia, Hungary, and Poland—but not to the Soviet Union. "In the American view, there's still a qualitative difference in the security situation between the Soviet Union and Eastern Europe," a senior administration official said.

Caere Corp. (Los Gatos, CA) claims that it has the **fastest optical page-reading system** on the market. The new Parallel Reader can process and recognize characters at a top speed of 700 characters per second, or 2500 words per minute, the company says. The $10,995 system's fast throughput is due mostly to its processing power and the way it handles pages. The box itself contains a 386SX CPU and four additional processor boards, each one holding a Motorola 68020 and 2 MB of RAM. Each 68020 board takes a page as it's scanned in and reads and processes the text. Other multiprocessor OCR systems assign different jobs to each processor and basically pass a page from one processor to the next, whereas Caere has all four boards work on different pages at the same time, said product support supervisor Dusty Perryman. The Parallel Reader comes with a 40-MB hard disk drive, a VGA display adapter, 1 MB of RAM, a 5¼-inch floppy disk drive, MS-DOS 4.1, and Windows 2.3; no keyboard, monitor, or mouse is included.

Tandy (Fort Worth, TX) has **slashed the price** of its low-end IBM-compatible 1000HX personal computer from $699 to $299. Introduced in 1987 as a home or school PC, the 1000HX comes with MS-DOS 2.11 and Tandy's DeskMate user interface in ROM. The computer supports an external 5¼-inch floppy disk drive (not included) and comes with Personal DeskMate 2.

situations in which the computer is literally moving while you use it (e.g., testing motor vehicles). The RM38 offers synchronous or asynchronous communications at up to 9600 bps. The 1-watt transceiver has a range of up to 10 miles, the company says.

Telesystems (Don Mills, Ontario, Canada) has taken a different approach than O'Neill: The Arlan 400 PC LAN uses transceivers that mount inside the computer (on an XT-bus board) and run a full version of Novell's NetWare. The data rate of the Arlan 400 is 230,000 bps—nearly 24 times the LAWN's 9600-bps speed—but the cards are also much more expensive ($1500 versus $495 per node). Telesystems president Wence Zenko says that his company is aiming for users who want to extend existing wired LANs between buildings. The radios have a range of 500 feet indoors and up to 6 miles outdoors across unobstructed space, Telesystems says.

Agilis (Mountain View, CA), maker of the eponymous "hand-held"

workstation, is selling its radio technology by the slice. The company's new 21CSP module ($2495) plugs into the Agilis computer and lets it act as a mobile workstation linked to an Ethernet LAN, with a data transfer speed of 230,000 bps. Each radio has a range of 1 kilometer outdoors and 100 meters indoors, Agilis says. The main bus on the Agilis works at an Ethernet speed of 10 Mbps. The transceiver connects directly to the Ethernet bus and converts the data rate to 236,000 bps for wireless transmission.

Agilis plans to offer its technology for other communications protocols, a company official said. The company recently showed us a prototype with a LocalTalk port on a 3- by 5-inch board. As it so happens, the transceiver's data transfer rate of 236,000 bps coincides with the transmission speed of the LocalTalk protocol used on Macintosh-based AppleShare and TOPS networks.

—Nick Baran and Andy Reinhardt

Mac OS to Get Print-to-Video Capabilities

By the end of the year, the Mac OS will incorporate video output capability, including the ability to print a file to videotape. Apple Computer is working with Sony to enable the Mac to write to videotape on Sony video machines that have serial control ports, according to Tyler Peppel, manager of Media Integration at Apple. Some of these videotape machines (VHS and 8-mm) retail for as little as $300.

Apple wants to integrate sound and video manipulation in the Mac OS so that all Mac programs can take advantage of such multimedia

capabilities, Peppel said at a recent conference on CD-ROM. Using a prototype of Mac System software, he was able to move a videotaped sequence to a single slide (created with the Power Point presentation program); when he clicked on a video display icon, the slide show ran with full animation and digital sound.

It would be as simple to write to videotape as to print to any other device, Peppel said. First you would select the videotape icon from the Chooser and then pick NTSC or PAL output and print.

—David Reed
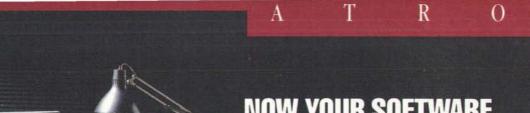
ISDN Puzzle Pieces Slowly Falling into Place

Skeptics are still willing to take bets: Which will come first— ISDN or human colonies on Mars? But some observers of the communications industry now say that the Integrated Services Digital Network— which promises high-speed voice and data (including video) transmission over a regular phone line—is finally "a reality" due to several important developments in the last few months.

Most significant, they say, is that

vendors and standardization groups have finally agreed on the Signal System 7 (SS7) standard for connecting islands of ISDN service to one another through conventional telephone switches. This will allow phone companies to connect the small pockets of ISDN service in cities such as Fresno and Chicago to one another, even though ISDN itself is not yet operating nationwide.

What the ex driving

Ask any PC enthusiast and they'll tell you that the Microsoft® Mouse is one very hot number.

Its nose is low and ominous. Its midsection rises up gracefully, connected by a smooth set of curves that flows across the entire body, creating a beautiful, sculpted look.

It's hard to believe this incredible form actually has a function.

But it does.

The patented design is made to fit the human hand. Making it easier to maneuver. And noticeably more comfortable. All of this advanced thinking didn't stop with looks.

Take a peek underneath the Mouse and you'll find the tracking ball in front, for better pointing accuracy.

With the new, expanded control panel there are hundreds of choices for sensitivity, pace and acceleration. Putting you in total control of the cursor's on-screen behavior. From meticulously slow right on up to warp speed.

perts will be this year.

And, with the resolution coming in at a mind-reading 400 points per inch, cornering, stopping and starting are second nature.

What it means is you can race through even the most complicated software programs without hesitation.

Combine all of this with the fact there's OS/2 support in the driver, and your Mouse is ready for today's sophis-

ticated applications.

You'll even be in control of what type of software you can add. Choose either Microsoft Windows™ version 3.0 or Microsoft Paintbrush™.

To get more information about the Mouse, or if you'd like to take a complete and thorough test drive, see your local Microsoft dealer.

One spin through your favorite program and you'll realize there's nothing else like it.

Microsoft®
Making it all make sense™

Texas Instruments (Dallas) **cut the price** of its TravelMate LT286 laptop computer. The Model 25, which comes with a VGA screen and a 20-MB hard disk drive, now costs $4599 instead of $4999. And the Model 45, which has a 40-MB hard disk drive, now costs $4999 instead of $5599.

If you buy seven of HP's top-line personal computers (386SX-, 386-, or i486-based Vectras), the company will throw in **a LaserJet III for free**. That's a $2395 value. Buy five, and HP throws in a LaserJet IIP. Buy three, and you get a free DeskJet Plus.

Irwin (Ann Arbor, MI) has **knocked $400 off the price** of its tape backup system for the Macintosh. The Model 5040, which is capable of storing 40 MB on a cartridge, is now $995. The Model 5080, which can back up 80 MB on a tape, is now $1295. Irwin recently started offering new tapes that will let both drives store an extra 20 MB per cartridge.

Advanced Micro Devices (Sunnyvale, CA) is **reducing the price** of the 16-MHz version of its 286 microprocessor by 30 percent. AMD says that demand for its 286 chips increased by 100 percent in the past year. The new price: $26 per chip in quantities of 1000, down from $37.

Cyrix (Richardson, TX) is now offering higher education a **discount on its FasMath coprocessor**, which is compatible with Intel's 80387. The new program lets a teacher buy two of the chips at half the normal price. The 16-MHz version normally sells for $570, the 20-MHz for $647, the 25-MHz for $814, and the 33-MHz for $994.

Astute observers noticed a bunch of ersatz reporters representing **"Computer Publicity News"** at IBM's RS/6000 rollout. They were, in fact, public relations employees from workstation maker Sun Microsystems.

Another major development is a final agreement on the so-called U-interface that governs the two-wire connections linking customer premises to local telephone lines. With these two pieces of the puzzle in place, ISDN is now standardized enough for large-scale implementation. One ISDN insider says that installation of large digital switches is proceeding faster than expected.

One thing holding back ISDN is the installation of fiber-optic cables all the way to the "local loop," or individual user site. Replacing existing copper wiring around the U.S. is expected to take another 5 to 10 years.

—Andy Reinhardt

IBM Promises Interoperability Between SAA, AIX

Although IBM has been heavily promoting its "seriousness" about Unix, manifested by its new RISC workstations and new version of AIX (IBM's rendition of Unix), there's a gap between IBM's Unix and its other operating systems. Those other operating systems—MVS, VM, OS/400, and OS/2—all comply with IBM's Systems Application Architecture. SAA is a set of guidelines for consistent interface and software development characteristics, designed to let SAA-based applications communicate and share data transparently.

The environments offered for the new RISC System/6000—AIX, OSF/Motif, AIX/Windows, and NextStep—do not conform to SAA.

IBM sees SAA as its main software architecture for its business customers, and AIX as the platform for its scientific and technical custom-ers. In other words, IBM expects that banks, insurance companies, accounting firms, and financial institutions will use mainly PS/2s, AS/400s, and mainframes, running OS/2, OS/400, or MVS and VM. Engineers, financial and statistical analysts, and government contractors will use mainly IBM's new workstations running AIX.

But rather than have those two classes of users isolated from each other, IBM officials say that they hope to alleviate the incompatibility between AIX and SAA by providing "interoperability" capabilities. Those capabilities will include support for sharing files and databases, common programming languages, network management, E-mail exchange, and support for the X Window System and Presentation Manager on both SAA and AIX platforms.

—Nick Baran

Imagine Buying an i486 for 35 Cents

By the year 2000, manufacturing semiconductors will have become so inexpensive that chips as complex as the i486 will cost about 35 cents (currently about $800), and memory chips with 1-MB capacity will sell for a dollar, predicts one industry watcher. Andy Rappaport, president of the Technology Research Group (Boston), said at the IEEE Compcon Spring '90 conference that if historic and current trends continue, the cost of making semiconductors will be minimal, and consumers as well as designers will benefit. Many electronic goods will become dramatically more powerful, Rappaport said, to the point that the "16-MB, 15-MIPS microwave oven" is a serious proposition. The cost of the semiconductors in a workstation capable of 30-MIPS operation will drop below $100 by the year 2000, he predicted.

—Owen Linderholm

PLUG US INTO YOUR NETWORK *and let us know what's going on. If you, your company, or your research group is working on a new technology or developing products that will significantly affect microcomputers and the way people work with them, we'd like to hear about it. Phone the BYTE news department at (603) 924-9281. Or send a fax to (603) 924-2550. Or write to us at One Phoenix Mill Lane, Peterborough, NH 03458. Or send E-mail to "microbytes" on BIX or to "BYTE" on MCI Mail. An electronic version of Microbytes, offering a wider variety of computer-related news on a daily basis, is available on BIX.*

LETTERS

The ASCII Controversy

Rip Collins' evaluation of ASCII (Stop Bit, January) is essentially correct. I never could comprehend the wisdom with which the punctuation and other symbols were arranged. The ASCII definitions with values greater than 127 seem tacked onto the end of the set as if by afterthought.

That's where my agreement ends. I don't understand how ASCII caused the divergence of compound delimiters in Ada, BASIC, and C. I agree that ASCII had something to do with the problem in not providing symbols such as "not equal to," forcing the framers of those languages to create their own. But I disagree that it forced everyone to produce different "not equal to" symbols. The choice seems arbitrary, depending on the framers' idea of what a comment or other symbol should intuitively look like.

Second, BASIC does not call text after a REM or an apostrophe "remarks" because comments are delimited by /* in C. BASIC was calling them remarks long before C existed. The issue is irrelevant to ASCII, anyway.

Third, programs do not "collapse under their own weight" on the basis of the order in which the language places the alphabet, digits, and other symbols. Well-written programs in those languages either use the order and layout of the character set and other system attributes to best advantage or find ways to circumvent the problems that those factors present.

Sure, ASCII order leaves something to be desired in some cases. Reordering the numerals and letters to facilitate character-to-hexadecimal conversion would be helpful in my line of work. However, pointing the finger at ASCII to such a degree seems unjustified.

Daniel M. Daly
Mishawaka, IN

Sorting the *interpretations* of unrelated terms would not be the phenomenal undertaking that the ASCII standard mandates were it not for the developers who promote "information interchange" as a sanctioned poetic license. Any Joe Blow who cares if "zany" precedes "aardvark" has never considered the computational poise that the ASCII standard provides in linking database, word pro-

TIME TO REPLACE ASCII?

Yesterday's character set is no match for today's computers

cessing, and spreadsheet functions with relative invisibility.

H. de Witt Robinson
Wilmington, DE

Apple Talks on AppleTalk

At Apple, we were pleased to read Mark L. Van Name and Bill Catchings' fine article entitled "AppleTalk Phase 2 and You" (January). Thanks.

Let me make one correction and provide clarification. The authors have equated an AppleTalk Phase 1 *zone* with a single network *segment*, and hence assert that zones are limited to 254 devices. This has never been the case. An AppleTalk zone is an arbitrary collection of network segments in an AppleTalk internetwork; although each segment is limited to 254 devices, a zone has no such limit. AppleTalk Phase 2 goes further. It allows any device on a given segment to choose the zone to which it wish-

es to belong. The network administrator can establish the list of zones from which the devices on a particular segment are allowed to make this decision.

On the subject of LocalTalk, it is clear that the authors have not examined the issue from the point of view of a customer's pocketbook. They say Apple has relegated LocalTalk to the low end. I suggest they avoid the potentially pejorative association of "low" and view the facts.

LocalTalk was conceived as and continues to be the lowest-cost LAN option today. We at Apple were not interested in engaging in "LAN locker room" contests ("my network goes further and faster"). Our interest was simply to allow the customer a penalty-free choice between LAN technologies at various price and performance points. LocalTalk has filled a real need.
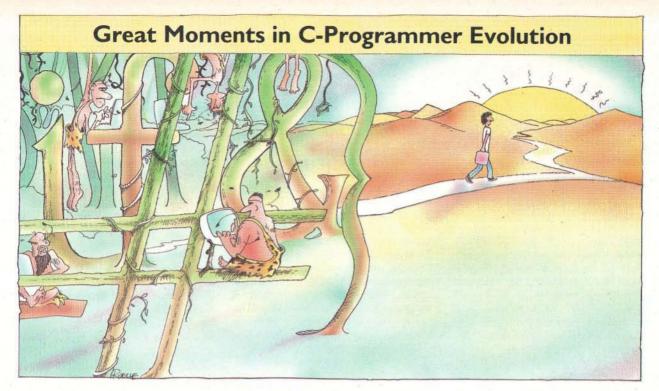
As for a faster LocalTalk network, the question is whether it is worth creating yet another option; the customers' vote on FlashTalk will help answer that. By the way, it is not the "only option," as you style it. Don't forget good old Ethernet and now also Token Ring.

The authors mention that now, with the TokenTalk NB product, more than one protocol stack can be run simultaneously. This is not something new in the AppleTalk system. The same has been true of LocalTalk and EtherTalk/Ethernet, as users of Apple's MacTCP and many other stacks will attest.

Gursharan Sidhu
Technical Director
Network Systems Development
Apple Computer, Inc.
Cupertino, CA

Thank you for your correction about zones. We differ, however, on your second major point; we regularly consider networking issues "from the point of view of a customer's pocketbook." We stand by our statement that Apple has relegated LocalTalk to the low end, with all its good (low-cost) and bad (low-performance) implications. For only a few extra dollars per Mac, Apple could easily increase the LocalTalk performance of at least the newer Macs to match that of FlashTalk, but to date Apple has chosen not to do so. We admire Apple's low-hassle, built-in networking support, but we still think

that LocalTalk is slower than desirable.

You also quote our comment on Flash-Talk as the "only option" out of context; it was, at the time we wrote that column, the only option that gave you higher performance over LocalTalk cabling. (DaynaTalk has since joined the fray.) Obviously, as the rest of our column dealt with them, we did not "forget good old Ethernet and now also Token Ring."

Finally, our statement that "More than one protocol stack can run on a single TokenTalk NB card simultaneously" does not imply that this ability is new to Apple networking; we were simply explaining how the TokenTalk NB card works.

We're glad to find Apple reading NetWorks, and we plan to cover more Mac networking issues in the future.
—Mark L. Van Name and Bill Catchings

DEC's Many Operating Systems
The assault from below has not forced Digital Equipment Corp. to abandon its "one architecture, one operating system" philosophy ("Computing with Light," January). DEC never had such a philosophy. It is conceivable that DEC management was sincere when it coined that slogan, but, if so, it never followed through. Even on the VAX, "one operating system" didn't last long. Ultrix is about as far from VMS as you can get.

Even within the DEC proprietary world, there was never a single architecture. The PDP-11 is still with us, and DEC sold a lot of them as Pro 350 and Pro 380 personal computers. The DECmate is a PDP-8, about as different from a VAX as is possible.

Seymour J. Metz
Annandale, VA

Uh-Oh, Call the Lawyers
Howard Eglowstein was absolutely correct in his response to Richard Levey's concerns about his hard disk drive write-protection modification (Ask BYTE, January). Not only is there no reason for concern over data errors or invalidating warranties, but hardware write-protection is the only "foolproof" way to protect a file on a PC hard disk.

I know that Eglowstein's modification works well, because I developed the same one four years ago and was issued a U.S. patent for it two years ago. Several hundred have been sold under the name Disk Defender, and customers have reported no problems. Great minds think alike, but if you want to give credit where credit is due, the patent office can document that I got there first.

Dennis Director
Evanston, IL

Mac Attacked
MacZealotry strikes again. I hope no one followed Don Crabb's advice by purchasing a Mac Plus with 1 megabyte of RAM and then tried to run applications under System 6 (Macinations, December 1989). He must have been kidding! System 6 itself barely runs in 1 MB.

He made another misleading statement: "This minimally acceptable OS/2 box will cost well over $3500." I can buy a 25-MHz 386 with 4 MB of RAM, a 65-MB hard disk drive, and color VGA video for under $2500 in any of several computer stores in the Seattle area. A Mac of similar performance and configuration would cost over twice as much. This is the beauty of the competition in the IBM PC–compatible marketplace.

Crabb seems confused about the difference between an operating system and a graphical user interface. I agree that OS/2's GUI, Presentation Manager, is not as refined in appearance as that of the Mac. However, OS/2 is considerably more refined than System 6.

Much of what Crabb calls refinement is merely user preference—he is used to the Mac way of doing things. I find some aspects of the Mac interface irritating; for example, all applications share one menu bar, inactive windows go blank, and you must hold down the mouse button while making a menu selection.

It is unfortunate that Crabb did not perform an unbiased comparison of the usefulness of OS/2, System 6, and their associated GUIs and applications, rather than making vague references about their level of "refinement." Such an approach would have been much more useful.

Waid Reynolds
Seattle, WA

Third-Party Apollo Vendor
In "Apollo Shrinks the Workstation Price Tag" (January), Ben Smith states that all options for the Apollo DN2500 are available only through Apollo. At Digital Micronics, our business and specialty is offering innovative solutions and options for upgrading Apollo workstations, including the DN2500. We offer memory boards, SCSI disk drives, tape drives, and high-resolution monitors.

Mark S. Bonney
Marketing Manager
Digital Micronics, Inc.
2131 Las Palmas Dr., Suite F
Carlsbad, CA 92009
(619) 931-8554

Helping Hands Update
We applaud Wayne Rash's column entitled "A Helping Hand" (January). How-

ever, the robotic arm that he mentioned was not designed by Prab; it was designed and is produced by the UMI Group Ltd., UMI-Microbot's parent company in the U.K.

The arm (which is designated by the RTX model number) is the most standard robotic system used in research by the leading medical/rehabilitation institutes. While the Prab Command is a sound generic approach, many other specific needs are being addressed by the medical research community, and there are many open areas needing work.

Interested readers can contact George Novelli, Managing Director, UMI Group Ltd., UMI House 9-15 St. James Rd., Surbiton, Surrey, KT6 4QN, U.K.
G. W. Rhodes
UMI-Microbot, Inc.
Sunnyvale, CA

All OS/2's Children
Mark J. Minasi's column on OS/2 multitasking (OS/2 Notebook, December 1989) uses the parent/child example to demonstrate the effects of delta and class on time slicing. This is fine as far as it goes, but the child programs fail to take advantage of one of the main features of OS/2's multitasking abilities—the use of shared code.

As he points out, the only difference between Child 1 and Child 2 is that they display 1 and 2, respectively, when reporting the number of iterations performed. So all that is required is a single Child program that passes its number as a parameter. Because OS/2 does not permit code segments to be written to, it can allow multiple processes to share the same in-memory copy of the .EXE file, each process having a separate copy of the data and stack segments.
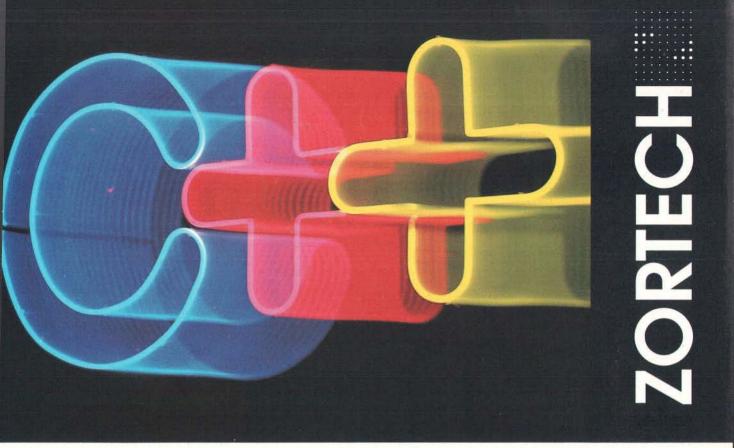
Bryan Ford
London, U.K.

The Primes of Our Lives
For a grizzled veteran of programming, it is amusing to see that prime number generation still enthuses and apparently puzzles younger programmers; one of whom, Milton Pope, reported that the Sieve of Eratosthenes runs twice as fast if you test only the odd numbers for primality (Letters, November 1989).

It is possible to increase the speed and reduce the memory requirements still further by a simple technique. This is based on the observations that the product of the first three primes is $2 \times 3 \times 5 = 30$, and that all other primes, when divided by 30, must produce one of the following eight remainders: 1, 7, 11, 19,
continued

23, or 29. Therefore, you can represent 30 consecutive numbers in 1 byte, so that an array of 32K bytes can represent a table of all the primes in 960,000 consecutive numbers. If you want to go further than this limit, you can allocate a new array block for each 960,000 numbers that you want to check. This can go on until you run out of memory, patience, or both.

If you want to get fancy, you use blocks of $7 \times 11 \times 13 \times 17 = 17,017$ bytes. When you have tested the prime numbers of 7, 11, 13, and 17, you copy the first block of 17,017 bytes, and in the copy you turn on the bits in the first byte that indicate divisibility by these four primes. Now you can use this modified copy to initialize all subsequent blocks of 17,017 bytes instead of testing them for divisibility with 7, 11, 13, and 17.

Søren Brandt
Rabylille, Stege, Denmark

Another Voice Heard

I was dismayed to discover that the December Resource Guide for the Sound and Image Processing In Depth section did not list our VPC-1000. It is an IBM PC–compatible option board that recognizes continuous speech, is speaker-independent, and operates over standard telephone lines. The VPC-1000 package sells for $5500 and is also available from OEM suppliers such as DEC, Summa Four, Syntellect, Perception Technology, and Periphonics.

C. M. Brooks
Executive Vice President
Voice Processing Corp.
One Main St.
Cambridge, MA 02142
(617) 494-0100

Don't Forget Number Crunching

Fred Langa's editorial in the December 1989 BYTE overlooked a very important concept: number crunching. The 80387 or 80387SX requires two to three times fewer clock pulses for a typical math operation than the 80287 does. For example, FMUL ST(i),ST(0), a multiply between two registers, takes 29 to 57 clock pulses on an 80387SX or 80387, and 90 to 145 clock pulses on an 80287.

If you have to do numerics processing, the 386SX/80387SX combination is far superior to the 286 or 80287.

Darrel L. Bartelheimer
Dayton, OH

Guidelines Needed for CASE

Your Product Focus entitled "Making a Case for CASE" (December 1989) was interesting, coming as it did after I had attended a brief methodology course at work. At my office, we have the KnowledgeWare IEW/WS set of tools, which I have found difficult to apply to my projects involving data transfer and processing across different platforms. Guidelines for determining which tools to apply to different types of problems are sorely needed.

Marc Walter
Loveland, OH

Big vs. Little

I enjoyed the review of the DECstation 3100 ("DEC's RISC Powerhouse," November 1989). I found the discussion about big- and little-endian machines interesting in a hardware review. I think this is primarily a software implementation issue. Network programs have the same problem that binary databases face. A native byte order needs to be defined, and the machines that don't conform to this native byte order need to flip bytes upon I/O. If the design takes this into account, moving to a different-endian machine merely requires recompiling the application. Addressing the issue at the hardware level reflects a lack of forethought regarding portability in designing software systems. If applications are intended to be portable, the byte order should not matter.

Marty Leisner
Rochester, NY

Enough, Already

I wonder if Pete Wilson ("The Wages of Sin," *IBM Special Edition*, Fall 1989) has any experience in the advertising and marketing fields. How many of us, I wonder, spend money every year on the latest and greatest version of a product simply because the manufacturer has convinced us that we don't know what we want and are not in control of our own needs? The demands that we place on hardware are more speed and compatibility with software that we already know, love, and have paid for.

Ric Naff
Dallas, TX

Who's Challenging Whom?

The November 1989 editorial ("What Slump?") states that the Motorola 68040 has not yet arrived to compete against the Intel i486. It seems to me that the i486 is now finally arriving to challenge the Motorola 68030, which has been in widespread use for over a year.

Let me refer you to BYTE's own benchmarks in the November 1987 issue ("MPW Compiler Lets Mac II's 68020 Shine"), comparing a Mac II with a 68020 to a Compaq Deskpro with a 386. Both ran at 16 MHz and had math co-processors. In five tests, the 68020 was faster in three (for the Dhrystone test only, higher numbers mean faster performance):

	68020	386
Dhrystone	3278	748
Fibonacci	48.58	53.11
Savage	3.42	8.95
Sieve	4.73	5.98
Sort	7.17	5.58

[All times are in seconds except for the Dhrystone, which is in Dhrystones per second.]

Furthermore, versions of the 68030 running at 40 to 50 MHz are now available. Given that you can generally expect commensurate changes in each family of CPUs as the clock speed changes, it is clear to me that the 68030 today is where the i486 will be in two years.

Mel Jenkins
Key West, FL

World's First E-Mail System?

Would anyone care to dispute my claim of having written and implemented the first E-mail system? I wrote it in 1975 for a DEC PDP-11/45 running RSTS/E 6.0. It allowed a user to send mail to another user—either in real time, preempting his or her current screen, or to a "mailbox" for future delivery. Any reader wishing to dispute this claim can write me at the address below.

David L. Morris
11311 North Central Expy.,
Suite 300
Dallas, TX 75243

ASK BYTE

Spelling It Out

I became permanently disabled in 1980 due to viral encephalitis, which resulted in a right temporal lobectomy. Upon my release from the hospital, I was given a used Apple II Plus computer and an ImageWriter printer. I am writing an autobiography, which will describe my illness, surgery, two-month coma, and several years of intensive rehabilitation.

I have learned how to use only one main program—Bank Street Writer. I am searching for a spelling checker program. I have been told that my computer, being an older model, does not have enough memory to use a spelling checker. Is there anything I can do at this

point? I would even consider trading my computer in for a newer model if that's possible.

Bill Hamilton
Bryan, TX

You don't say how much memory your computer has, so I'm going to assume that it's 64K bytes. If not, you can probably find 16K-byte upgrades for a reasonable price at any number of mail-order houses. But 64K bytes should be sufficient for what you're planning to do.

Sensible Speller IV from Sensible Software (335 East Big Beaver, Suite 207, Troy, MI 48083) is compatible with Bank Street Writer. Again, search for it from a mail-order house; a company like Programs Plus (75 Research Dr., Stratford, CT 06497, (203) 378-3662) should carry it. Be sure to ask for Sensible Speller IV—it will work on your Apple II Plus. Sensible Speller without the "IV" will not.

—R. G.

An AT in XT's Clothing

About a year ago, I purchased a Tandy 1000 TX. At that time, the salesperson gave me the impression that the TX was an AT-class machine. When the time came to upgrade the computer, I found out that it had an XT motherboard with a 286 chip for a processor. This really ticked me off, particularly because it was my fault for not reading the fine print.

My question is: Do I have an XT or an AT machine? I haven't been able to get a straight answer from anyone.

Another concern that I have is buying additional components for the computer. I would like to buy these options through mail-order companies in the U.S. (due to the high price of these components offered by Tandy here in Nova Scotia). However, people I talk to here give me opposing opinions as to whether my machine will accept printers, modems, and graphics cards manufactured by other companies. Do I have to buy components from Tandy?

James Erwin
Stellarton, Nova Scotia, Canada

You have an AT-class computer with an XT bus. The 1000 TX was Tandy's attempt to couple the power of a 286 with the economy of using 8-bit expansion cards. You should be able to use any 8-bit card that measures 10 inches or less in your computer and connect your 1000 TX to any standard printer or modem. Some advertisers in BYTE sell peripherals (e.g., hard disk drive cards) specifically for Tandy computers.

You cannot upgrade the graphics on

the 1000 TX, since it uses a built-in CGA circuit on the motherboard that you can't bypass.—S. W.

Speaking the Language

I use Windows/386 extensively with my Compaq 386s. My monitor is the original Compaq VGA color monitor. One software package that I often use is Samna's Ami.

Here in Brazil we speak Portuguese, and the Portuguese language has some characters that are not found in English. I used to access such Portuguese-unique characters through the Alt keys on my Brazilian version of Word 3.0. (This version of Word has Alt keys adapted to special Portuguese characters.)

Enter Windows/386. Using Ami—or even Write (which comes with Windows/386)—not all the special characters are available. I understand that Windows/386 is a graphical environment. I suppose there must be a way to modify the font files for these characters so that I could print in Portuguese. Am I right? If so, how should I proceed?

Carlos Roberto Fortner
Sao Paulo, Brazil

Portuguese? Some of us here consider C a foreign language, so I had to go right to the source.

Samna has no immediate plans to release a Portuguese-language version of Ami (or Ami Professional) in the immediate future, but a spokesperson pointed out that any version of Ami can access the full ANSI Windows character set. The Alt key, used with the numeric keypad, will let you enter the numeric code of any character. A table in the back of your manual will show you the available characters. Samna was very helpful, and I'm sure that if you give the company a call at (404) 256-2272, someone will be happy to work with you.

It is possible to modify the Windows font. The font file itself is essentially a table of bit maps, one for each character. The format is documented in several books and the Windows' SDK (Software Development Kit) documentation. Unfortunately, because the printer drivers do not use the same font file, hacking the font file won't help you, either. What you need is a version of Windows with a full Portuguese character set. For that, I called Microsoft.

Strike three. The problem, apparently, is that Portuguese has two characters that are not normally available in the ANSI set. A Portuguese version would have the following two things: (1) support for these two characters and (2) Portu-

guese menu prompts.

Microsoft is not yet shipping a Portuguese version of Windows, but one is currently in development. There's no word on when it will be released. In your area, you might call someone at Microsoft Brazil in Sao Paulo (55-11-5304455) and see if the people there can help you track it down.—H. E.

The Vectra and VGA

In the March issue, Frederik Wessels wrote about a compatibility problem between a Hewlett-Packard Vectra computer and a VGA video board. I may be able to help him. In my job, I take care of about 70 Hewlett-Packard Vectras. I have found a couple of things that make a Vectra work better with VGA cards.

Run the setup program that came with your computer. Under the "set system configuration" option on the main menu is a way to indicate what type of video card you have installed. Select VGA.

If that doesn't work, then you might want to go back to the setup main menu and enter the word YADA. That's right. The elusive "HP Support Personnel" menu will greet you. If you see a selection called "Shadow RAM video BIOS," you may want to turn it off. Your VGA card may not like to have its ROM BIOS shadowed into RAM, so disabling the shadow option may work.

Anyway, if neither of these suggestions works, then the problem is probably a conflict between the video card and Hewlett-Packard's ROM BIOS.

Steve Riley
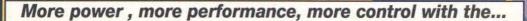*Computer Systems Engineer*
Ashland Chemical
Dublin, OH

Thanks!—Lab Staff

FIXES

• There are two corrections to "The Mac State of Mind" by Daniel Rasmus (January). First, Apple has no plans to develop or distribute Pearl Lisp. Second, "Macintosh Allegro Common Lisp" is the correct name for Apple's Lisp product.
• The 800 number for Enpower Corp., published in What's New in March, is no longer correct. Further information regarding Enpower will be made available if it comes to our attention.
• In listing 2 of "Flirting with Assembly" by Hugh Kenner (April), the fifth line should read as follows:
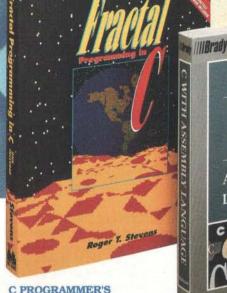$74/43/ { JZ L2 (+43) } ∎

Ultra-Comp's 386

The Ultra-Max 386-33 is an inexpensive system from Ultra-Comp that comes with a 64K-byte cache, 4 MB of RAM, a 160-MB 28-ms ESDI hard disk drive, VGA graphics (including the monitor), both 3½-inch 1.44-MB and 5¼-inch 1.2-MB floppy disk drives, and DOS.

The Micronics motherboard has a Phoenix BIOS, an Intel 82385 caching controller, and a math coprocessor port. A 32-bit memory slot handles up to 8 MB (up to 16 MB with a daughterboard). There's also room for five 16-bit and two 8-bit cards.

The keyboard has 101 keys, and the 8-bit I/O board has one parallel and two serial ports. The VGA graphics is controlled by a 16-bit Genoa board with 512K bytes of RAM and displayed on a 14-inch Sony Super VGA monitor with a resolution of 1024 by 768 pixels.
Price: $4595; add $150 for tower configuration.
Contact: Ultra-Comp, 11988 Dorsett Rd., Maryland Heights, MO 63043, (800) 435-2266 or (314) 991-1988.
Inquiry 1120.

Tower or desktop, the Ultra-Max 386-33 is packed with features.

ALR Shrinks Its Tower

The FlexCache 25386DT Model 40 from ALR is a redesign of the company's 386/25 tower system. It includes 64K bytes of cache memory, 1 MB of RAM, a 5¼-inch 1.2-MB floppy disk drive, a 40-MB 28-ms hard disk drive, a 101-key keyboard, and a 200-W power supply.

There's also room to expand with three free half-height bays and one 8-bit and six 16-bit slots. The hard disk drive controller is located on the motherboard, and up to 16 MB of RAM resides within the 32-bit memory slot. The floppy disk drive controller and serial and parallel ports are on one 16-bit card, which takes up one of the slots.

The FlexCache Model 40 measures 21½ by 17 by 6 inches and weighs 60 pounds.
Price: $2695.
Contact: Advanced Logic Research, Inc., 9401 Jeronimo, Irvine, CA 92718, (714) 581-6770.
Inquiry 1123.

Battery Not Included

Grid Systems' new 8½-pound backlit-VGA laptop doesn't come with a battery, but the optional nickel-cadmium battery pack attaches to the bottom of the system and looks like a base. The 2-hour battery weighs 3 pounds, and the 4-hour battery weighs 4 pounds.

The Grid 1450sx comes with a 3½-inch 1.44-MB floppy disk drive, a 2½-inch 20-MB hard disk drive, a 10-inch diagonal VGA monitor, an 82-key keyboard, 1 MB of RAM, and a Phoenix BIOS. Standard ports include two RS-232C serial, one Centronics parallel, and one connector for an external monitor. You can expand RAM to 5 MB, and a 2400-bps internal modem with MNP 5 data compression is available.
Price: $4995; 2-hour battery, $299; 4-hour battery, $399; modem, $595; cigarette-lighter adapter, $25.
Contact: Grid Systems Corp., 47211 Lakeview Blvd., P.O. Box 5003, Fremont, CA 94537, (800) 222-4743 or (415) 656-4700.
Inquiry 1121.

Casio Introduces Fourth B.O.S.S.
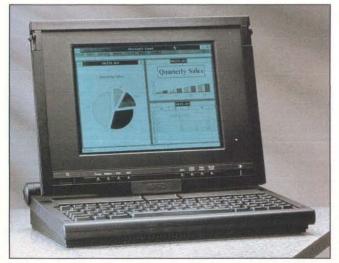
The SF-9000 Executive B.O.S.S. (Business Organizer Scheduling System), a new notebook computer from Casio, uses four types of expansion cards.

The SF-9000 comes with 64K bytes of RAM (expandable to 128K bytes with the ES-100 memory card); a 6-row by 32-column display; a separate business card library and telephone directory; a memo function; a side-by-side two-month calendar display, schedule, and daily alarms; and world-time functions.

The SF-9000, which has a QWERTY keyboard, measures ¾ by 3¼ by 6 inches and weighs 8 7/10 ounces. Three cards are optional: the ES-600 electronic dictionary, the ES-610 financial/legal spelling checker, and the ES-620 medical spelling checker.
Price: $299.95; ES-100, $129.95; ES-600, $129.95; ES-610, $79.95; ES-620, $99.95.
Contact: Casio, Inc., 570 Mt. Pleasant Ave., P.O. Box 7000, Dover, NJ 07801, (201) 361-5400.
Inquiry 1124.

Grid Systems' 1450sx laptop weighs 8½ pounds without a battery.

A Rewritable Optical Drive for DOS, Mac, or OS/2

The LaserVault 650 is a rewritable optical subsystem that lets you store data generated from within DOS and the Mac OS. In July it will also support OS/2.

The drive is a 95-ms Sony SMO-S501. Tecmar includes a SCSI board (ISA or Micro Channel architecture) and Tecmar's custom software.

The key to LaserVault's multiple operating-system compatibility is the single-partition format, Tecmar says. You use it to partition sections for your DOS files, sections for your Mac files, and sections for your OS/2 files. You can make partitions almost any size. And a feature Tecmar calls "DOS compatibility" lets you mix DOS 2.x through DOS 4.x files. "DOS compatibility" also eliminates the 32-MB partition limit for MS-DOS versions earlier than 3.30.

Software features include a manager function for setup, partitioning, and diagnostics; a device driver for direct access; and a backup manager for archiving.

The setup program automatically formats and partitions new media and lets you set up your own partitioning requirements. Diagnostics perform drive and media testing. And the archiving portion of the software supports automated backup.
Price: $5495; interface kits, $200 to $600.
Contact: Tecmar, 6225 Cochran Rd., Solon, OH 44139, (216) 349-0600.
Inquiry 1125.

The LaserVault 650 is a rewritable optical drive with multiplatform SCSI drivers.

Low-Priced Color Monitor Ships with VGA Controller

The VisionMaster PMV14VC/plus is a 14-inch color VGA monitor for your PC or Mac that features a dot pitch of 0.28 mm and a maximum resolution of 1024 by 768 pixels.

The interlaced or noninterlaced monitor comes with a three-quarter-length VGA (640- by 480-pixel) card that has 512K bytes of RAM. Other features include a bandwidth of 45 MHz, a vertical scan frequency of 50 to 90 Hz, a tilt-and-swivel base, a nonglare screen coating, brightness and contrast controls, and a horizontal scan range from 31.5 to 35.5 kHz.
Price: $789.
Contact: Sirex, Inc., 132-14 11th Ave., College Point, NY 11356, (800) 722-0404 or (718) 746-7500.
Inquiry 1127.

Store Files on 700-MB DAT

MegaDat is an inexpensive 700-MB 4-mm digital audio tape storage system for PCs that features a 60-second maximum access time.

It supports both sequential and direct access, is front-loaded, and can be configured as an internal or external peripheral. You need NetWare 286, DOS 3.0, 1 MB of RAM, a hard disk drive, and a 16-bit slot.
Price: Internal, $2750; external $2950; one tape, $36.
Contact: GigaTrend, Inc., 2234 Rutherford Rd., P.O. Box 4298, Carlsbad, CA 92008, (619) 931-9122.
Inquiry 1129.

SPREAD THE WORD
Your new product is important to us. Please address information to New Products Editors, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458. Better yet, use your modem and mail new product information to the microbytes.hw or microbytes.sw conferences on BIX. Please send the product description, price, ship date, and an address and telephone number where readers can get more information.

Scan Documents with the Xerox Datacopy GS plus

Xerox offers its 8-bit flatbed scanner with software for your Mac, DOS, or OS/2 system and a package that includes desk accessory software for Windows-compatible desktop publishing with Presentation Manager, PageMaker, and Ventura Publisher.

The Datacopy GS plus has a SCSI port for your Mac and an 8-bit SCSI add-in card for ISA and Micro Channel–based systems. The scanner is adjustable for start-stop times, defining read areas, and contrast.

It will scan documents up to legal size with scanning densities of between 75 and 300 dpi in 1-dpi increments. Typical scanning speed is 21 seconds for a two-shade letter-size document. Output modes include two-level, 256-level dither, and 4- or 8-bit gray scale.

The Datacopy GS plus weighs 22 pounds and measures 21½ by 13⅞ by 4⅞₀ inches. Options include the Automatic Document Feeder.
Price: $2195; document feeder, $595.
Contact: Xerox Imaging Systems, 535 Oakmead Pkwy., Sunnyvale, CA 94086, (800) 248-6550.
Inquiry 1128.

continued

Speed AutoCAD Redraws by Factor of 25

The Artist XJS is a new 34020-based graphics co-processor board designed for high-end CAD applications. Coupled with an Artist GT software driver, which uses display-list processing, the XJS board performs AutoCAD redraws at speeds of up to 25 times faster than if you used the standard AutoCAD driver, according to the manufacturer.
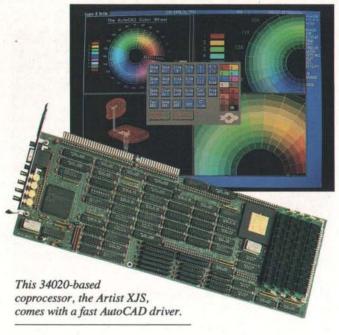
The base controller has 1 MB of video RAM and offers 1280- by 1024-pixel resolution with 16 colors or 1024- by 768-pixel resolution with 256 colors. Add another megabyte of video RAM and you get 1280- by 1024-pixel resolution with 256 colors and 1600- by 1200-pixel resolution with 16 colors. A palette of 16.7 million colors is standard with all versions.

The base controller also has 1 MB of DRAM, which stores display lists. You can upgrade the DRAM with an optional 4 MB of user-installable single in-line memory modules.

You can use the XJS board in single-screen and dual-screen systems. You can use single-screen VGA pass-through using a multiscanning monitor.
Price: $3495; with 2-MB daughterboard, $4295; separate user-installable 2-MB daughterboard, $1000; VGA daughterboard, $795; VGA link, $249.
Contact: Artist Graphics, 2675 Patton Rd., St. Paul, MN 55113, (612) 631-7800.
Inquiry 1177.

This 34020-based coprocessor, the Artist XJS, comes with a fast AutoCAD driver.

IBM to Offer PS/2 Board and Software for Multimedia Work

IBM's new M-Motion Video Adapter/A is a Micro Channel card designed for multimedia. You use it to digitize and display video signals, and you can digitize sound and play it through a speaker.

The M-Motion Video Adapter/A can take National Television System Committee or phase alternate line input from video cameras, VCRs, videodisks, and TVs. You can manipulate motion or still images within windows, change the size of the windows, and lay text or graphics over the video images. You can use up to two standard stereo sources.

The software, the M-Control Program, provides tools for developing multimedia applications with the M-Motion Video Adapter. The package lets you control video and audio operations, such as changing the sharpness and color of an image on the screen or choosing the source of the input.
Price: $2250; M-Control Program, $150.
Contact: IBM, U.S. Marketing and Services, Westchester Ave., White Plains, NY 10604.
Inquiry 1134.

Add RAM in SIMMs to Your Deskpro 386s

Amkly Systems offers memory boards for your Compaq Deskpro that let you add 32-bit RAM modules for less than the competition.

The AMpac series boards all fit the Deskpro's 32-bit memory slot and let you upgrade your Deskpro's standard 2 MB with up to 13 MB of single in-line memory modules. The AMpac-SIMM-1000s and AMpac-SIMM-4000s feature 1 MB and 4 MB, respectively.
Price: AMpac-SIMM1000s, $695; AMpac-SIMM4000s, $1795.
Contact: Amkly Systems, Inc., 60 Technology Dr., Irvine, CA 92718, (714) 727-0788.
Inquiry 1133.

RasterOps Board Puts Live Video on Mac Screen

With RasterOps' Nu-Bus ColorBoard 364, you can inexpensively create full-motion video with your Mac. It captures video in 24-bit color from any National Television System Committee (NTSC) source—such as a TV, a VCR, or an S-Video recorder—and can display it in real time (30 frames per second) on any noninterlaced Macintosh screen.

Your Mac II can then display live, incoming video in a resizable window while running another software application. RasterOps claims that the ColorBoard 364 can run the live video simultaneously with any application without any performance degradation or interruption. MultiFinder is not required.

Accompanying software lets you select pixel depths (to adjust color), virtual screen size (from 640 by 480 pixels up to 4096 by 512 pixels), hardware pan and zoom, and other options. The frame-grabbing application "grabs" frames in real time ($\frac{1}{30}$ second). A series of HyperCard XCMDs (external links to other applications or peripherals) are included for developing multimedia applications.

RasterOps also plans a series of ColorBoard 364 add-on products for conversion to NTSC or to receive and tune broadcast signals.
Price: $1995.
Contact: RasterOps Corp., 2500 Walsh Ave., Santa Clara, CA 95051, (408) 562-4200.
Inquiry 1132.

Share Your LaserJet III

With the new Share-Spool cards from Extended Systems, you can connect up to four computers to one LaserJet. You install a card in the I/O slot in the printer.

The higher-performance 2094A connects through optionally available custom twisted-pair cables and Centronics parallel ports. The 2041C uses serial ports.

With the 2094A, PCs can be located up to 100 feet from the LaserJet Series III, and it provides 1 MB of RAM for spooling. Each 2094A accepts up to four simultaneous data signals at rates approaching 240,000 bps each.

The 2041C includes 256K bytes of RAM (a 1-MB version is also available). Data rates approach 19,200 bps.
Price: 2094A, $845; 25-foot cable, $70; 100-foot cable, $140; 2041C, $495; 1-MB 2041C, $690.
Contact: Extended Systems, 6062 Morris Hill Lane, Boise, ID 83704, (208) 322-7163.
Inquiry 1135.

ShareSpool cards connect your LaserJet to as many as four computers. They also provide RAM for spooling.

No More Bad RAM

Simcheck is a portable memory module tester for single in-line memory modules and single in-line

packages.

It features a 16-bit 10-MHz V20 processor and an automated digital speed synthesizer for testing access time down to 20 ns. The two-line LCD displays module type, size, and speed. Simcheck then tests all data bits so you don't have to manually switch from bit to bit.

Auto loop test is a standard feature; Simcheck automatically generates test patterns. Simcheck also comes with programmable voltage sources for advanced testing techniques, including voltage bounce and voltage cycling. A chip-heat mode warms modules for temperature-dependent tests.

Innoventions says that Simcheck will identify defective chips within modules and open- and short-wiring problems and will provide repair information. Simcheck measures 5 by 7 by 1½ inch-

es and weighs less than 2 pounds.
Price: $995.
Contact: Innoventions, Inc., 11000 Stancliff, Suite 150, Houston, TX 77099, (713) 879-6226.
Inquiry 1137.

Add Memory to Your IBM LaserPrinter

LP Memory is a memory module that lets you add 1, 2, or 3.5 MB of RAM to your IBM LaserPrinter 4019.

Pacific Data Products says that upgrading memory will let you take advantage of IBM's memory manager, which dynamically allocates RAM for soft fonts, bit maps, or a spooling buffer, depending on what type of print job you're performing.

Additional memory also

lets you download fonts and forms, print at the maximum resolution of 300 dpi in desktop publishing applications, and perform higher-resolution printing in the plotter mode. You can also use LP Memory as a spooling buffer for the IBM printer-sharing option (for up to six PCs). The 4019 printer-sharing option is not buffered, relying on either the printer or processor memory for buffering activities.
Price: 1 MB, $395; 2 MB, $595; 3.5 MB, $895.
Contact: Pacific Data Products, 6404 Nancy Ridge Dr., San Diego, CA 92121, (619) 552-0880.
Inquiry 1136.

Control Stepper Motors from Your Parallel Port

A new kit and software called Indexer LPT from Ability Systems lets you control an included stepper motor through your computer's parallel port.

The starter kit includes a motor, a driver, a single-axis cable (for one motor), limit switches, and a manual that tells you how to use the Indexer LPT software. With the kit, you can control a motor axis with respect to position, home position, maximum and minimum velocity, acceleration, and automatic return to home.

LPT software, which loads as a device driver, can control up to six motors at once (through three printer ports). Included diagnostic software lets you troubleshoot motor problems.
Price: $600; software alone, $199.95.
Contact: Ability Systems Corp., 1422 Arnold Ave., Roslyn, PA 19001, (215) 657-4338.
Inquiry 1140.

Travel Overseas Without Worry About Foreign AC

The lapAdapt plug converter lets you use your laptop computer in many foreign countries without worrying that the AC will harm your machine or, more important, the data within, the manufacturer claims.

The grounded plug adapter converts 110-V to 240-V

AC to work with your American-compatible laptop's AC power adapter and adds a grounding plug. You plug your AC adapter into lapAdapt's grounded connector.

The grounded connector plugs into the British adapter plug for the U.K. (about 2 inches square on each end

with a 1-foot power cable in between) or into an included plug for other foreign outlets (also about 2 inches square).
Price: $49.95.
Contact: lapAdapt Electronics, 169 11th St., P.O. Box 611, Hoboken, NJ 07030, (201) 798-3030.
Inquiry 1138.

BUY A RACEHORSE...
GET A WORKHORSE!

FAST!
850cps/240 lpm!
And It Runs . . .
And Runs . . . And Runs . . .

The fastest serial dot matrix
printer on the market today!
The all new **850XL** offers a world of benefits!

- Lightning fast at 850 cps (240 lpm throughput)
- Continuous printing capabilities with no overheating
  or unnecessary downtime!
- Over 300 local service centers nationwide to keep
  your jobs running day and night!*

The waiting game is over, as the **850XL** takes
on mountains of data, round the clock, with no duty
cycle restrictions! Any printing application you need
is handled with rapid-fire reliability:

- Data Processing
- Financials
- Bar Codes
- Spreadsheets
- Near-letter Quality
- Labels
- Graphics

*Call for availability in your area.

Standard features are better than ever!

- 5 to 18.2 Pitch Printing
- Front Panel Menu
  Programming
  (No DIP Switches)
- Quietized Enclosure
- EPSON, DEC, and IBM
  ProPrinter XL Emulations
- 8K Data Buffers
- Serial & Parallel Ports
- Convenient Front &
  Bottom Paper Feed
- Full International
  Character Set

OTC . . . An American Winner!
Call today for more details.

1-800-4-OUTPUT (8 am - 5 pm PST)
(468-8788)

Call me, I'm interested: Circle 184

Pocket Modem Needs No Battery

The Stowaway 2400 is a pocket-size 2400-bps modem that uses telephone line power instead of a 9-V battery. Features include backward compatibility, auto-dial, auto-answer, and a 40-character command buffer.

All phone companies provide line power for tones and ringers at a guaranteed 20 milliamperes, and that's enough to power your Stowaway 2400 even if your telephone line is 20 miles from the nearest phone company generator.

The Stowaway measures 3 by ¾ by 2 inches and weighs 2⅛ ounces.
Price: $295.
Contact: VoCal Technologies, Ltd., 3032 Scott Blvd., Santa Clara, CA 95054, (408) 980-5181.
Inquiry 1141.
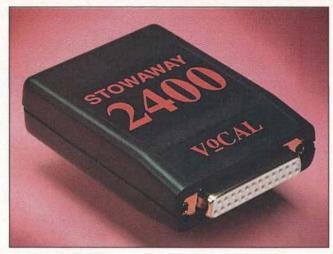
Low-Priced LAN Uses Parallel Ports

The PC-InterLink is a peripheral-sharing and file transfer network with a peer-to-peer network operating system, and hardware.

With the starter kit you can link two DOS-based personal computers. The hardware includes two parallel-port adapters, two 10-foot cables with RJ-11 connectors on the ends, and a four-port junction box. You can locate as many as four computers up to 100 feet from the junction box. The data transfer rate approaches 9600 bps, depending on your system's processor.
Price: Starter kit, $199.95; additional nodes, $49.95.
Contact: SoftWorx, 801 East Campbell Rd., Suite 355, Richardson, TX 75081, (214) 480-8278.
Inquiry 1150.

VoCal makes the Stowaway 2400 smaller by using telephone line power.

New Crosstalk Communications for NetWare LANs

New versions of R2LAN and Crosstalk XVI software now run on NetWare. R2LAN 2.0, a server-based remote-control program, now supports Novell's IPX/SPX transport protocols, as well as those of NetBIOS. Another new feature is the Guard Host Security System, which lets you control access to files on an individual basis. File transfer has also been enhanced, and hosts are now able to initiate calls to clients.

Version 3.71 of the network server version of Crosstalk XVI now has support for Novell's NetWare Asynchronous Services Interface, allowing you to communicate with a Novell LAN Asynchronous Communications Server.

Other features new to the network version include call logging, a command-line editor, support for files in other directories, and enhancements to Kermit and the program's scripting language.
Price: R2LAN, $795; Crosstalk XVI, $600.
Contact: Crosstalk Communications, 1000 Holcomb Woods Pkwy., Roswell, GA 30076, (404) 998-3998.
Inquiry 1145.

OS/2 Multitasking Communications Called MLink

MLink, a multiplatform data communications software package, has been redesigned from its DOS, Mac, and Unix roots to work with OS/2. It provides terminal emulation, error-free file transfer, data compression, and a built-in scripting language for automated operation.

It operates in OS/2's protected mode and can exploit such multitasking features as simultaneous communication from multiple ports, and interprocess communications. A development system version with a scripting compiler is also available.

A companion product, MLink Polling Manager, uses OS/2's multitasking as it works to automate file transfers with remote systems.
Price: MLink, $310; MLink Development System, $395; MLink Polling Manager, $695 to $3995.
Contact: Corporate Microsystems, Inc., Mt. Support Rd., P.O. Box 2059, Lebanon, NH 03766, (603) 448-5193.
Inquiry 1144.

Save RAM with LANsmart 2.0

LANsmart 2.0 is a peer-to-peer network operating system that can run workstations and servers in as little as 10K and 64K bytes of RAM, respectively. It can also run concurrently with NetWare 2.15.

The RAM conservation is due to LANsmart's use of about 64K bytes of extended memory on workstations and about 70K bytes of extended memory in your server.

You can purchase LAN-smart 2.0 in a starter kit that includes two ARCnet cards plus cabling (you must buy additional ARCnet cards separately). Ethernet and Token Ring kits are also available, as are cards.

If you want to run LAN-smart on top of NetWare, you'll need a separate product called LANsmart NW, which you load on your NetWare 2.15 server. One possible application for running concurrent operating systems is using one tape back-up system for the server and, through the peer-to-peer LANsmart, using another tape backup system for server-based tape backup of remote hard disk drives.
Price: ARCnet starter kit, $595; additional cards, $195; Ethernet or Token Ring starter kit, $1300; additional cards, $369; LAN-smart NW, $595.
Contact: D-Link Systems, Inc., 5 Musick, Irvine, CA 92718, (714) 455-1688.
Inquiry 1143.

Buy a portable and

T3200SX: 170 pounds, 16MHz 386SX with 80387SX-16 coprocessor socket, 5 built-in ports, 40MB hard disk with 25msec access, 1MB RAM expandable to 13MB, gas plasma VGA display with 16 gray scales, 1.44MB 3½" diskette drive. IBM is a registered trademark of International Business Machines Corp. 386 is a trademark of Intel Corporation.

For more information call 1-800-457-7777.

Two IBM-compatible expansion slots (one full-length and one half-length) and an internal modem slot let you connect to LANs, mainframes and more. Plus, three dedicated internal expansion slots let you increase memory up to 13MB.

TOPS Network Bundles for DOS and Mac

The Network Bundle for DOS and the Network Bundle for Macintosh include new versions of the TOPS file-sharing system and the new InBox 3.0 E-mail system. The Network Bundles are designed to provide the three most common network services—file sharing, printer sharing, and E-mail.

TOPS Network Bundle for DOS has the new features of TOPS 3.0, including extended memory support. TOPS occupies between 65K and 230K bytes of conventional memory, depending on the configuration of the system. Also, you now can configure the file- and print-sharing systems to include only those features you actually use.

TOPS users now can print DOS applications directly to networked Apple Laser-Writer printers. Also, because TOPS supports the Apple Filing Protocol (AFP) and the Hierarchical File System, a PC can operate as a file server for

File sharing and E-mail are now bundled by TOPS.

Macintoshes on your network.

TOPS Network Bundle for Macintosh works with AFP-compatible multiuser applications such as 4th Dimension, FoxBase+, and FileMaker II. It now supports the Apple-Talk limit of 254 zones and up to 254 servers per zone, and the new Apple LaserWriter

drivers (6.0). A new version of MacLinkPlus/TOPS from DataViz is included, replacing the file translators found in previous versions of TOPS/Macintosh.
Price: Network Bundle for DOS, $249; Network Bundle for Macintosh, $299; DOS or Mac site upgrade, $125, or

$185 for both.
Contact: TOPS, 950 Marina Village Pkwy., Alameda, CA 94501, (415) 769-9669.
Inquiry 1142.

NetWare Printing Made Easier

Two new versions of PS-Print, 3.0 and 3.1, are now available to help you print from any personal computer on your NetWare 2.1 to 2.15 LAN and NetWare 386 LAN, respectively.

PS-Print 3.0 gives you better flexibility in assigning print servers and is now compatible with NetWare-networked PCs that use noninterrupt-driven printer ports, Brightwork Development says. PS-Print 3.0 and 3.1 use about 8K bytes per printer on the print server.
Price: Site license for PS-Print 3.0 or 3.1, $395.
Contact: Brightwork Development, Inc., 766 Shrewsbury Ave., Jerral Center West, Tinton Falls, NJ 07724, (800) 552-9876 or (201) 530-0440.
Inquiry 1149.

A Trio of Fax Modems

The 2400PF Fax/Modem from Micro Electronic Technologies is a low-priced pocket-size V.22bis (2400-bps) modem coupled with a send-only V.29 (9600-bps, Group 3) fax device. The unit measures 1 by 4½ by 2¾ inches, including the 9-V battery.

The backward-compatible (to 1200, 600, and 300 bps) data communications device doesn't include the "pass-through" port of its recently introduced 1½- by 5½- by 4-inch 2400XF sibling but does have software for timing faxes and for personal phone directories.

Price: 2400PF Fax/Modem, $279; 2400XF, $239.
Contact: Micro Electronic Technologies, Inc., 35 South St., Hopkinton, MA 01748, (508) 435-9057.
Inquiry 1146.

If you're on the road and need to receive faxes as well as send them, then The Complete Fax Portable might serve you well.

It's a V.29 (9600-bps, Group III) send-and-receive fax device. The Complete Fax Portable measures 1 by 5¾ by 3 inches and weighs 10 ounces, including the 9-V battery that is rated for 3

hours of continuous use.

An included TSR program resides in about 60K bytes of main memory and features the usual modem/fax features: timed faxing, a database of frequently faxed numbers, direct scanning, and broadcast.
Price: $499.
Contact: The Complete PC, 1983 Concourse Dr., San Jose, CA 95131, (800) 634-5558 or (408) 434-0145.
Inquiry 1147.

The FaxDat from Xecom is a V.22bis/V.29 fax modem with send and receive capabilities in an add-

in card for your Compaq SLT/286, NEC ProSpeed 286, Toshiba T1600, or Zenith SupersPort 286.

When faxing, the card has fallback rates of 7200, 4800, and 2400 bps. If you use it for data communications, it includes MNP Level 5 error correction and data-compression technology and is backward-compatible with the V.22 and Bell standards. Each card is shipped with fax software.
Price: $625.
Contact: Xecom, Inc., 374 Turquoise St., Milpitas, CA 95035, (408) 945-6640.
Inquiry 1148.

Get the latest Word from SCO.

Microsoft Word 5.0 for UNIX Systems.

You've come to depend on SCO™ for the latest UNIX® System software solutions for PCs. Industry standards such as SCO™ XENIX® 386 and SCO UNIX System V/386 Release 3.2. World-famous applications such as SCO Professional® the 1-2-3® workalike, and SCO™ FoxBASE+™

And now Microsoft® Word 5.0, the same full-featured word processing system that has defined power, speed and flexibility for MS-DOS® and OS/2™ users, is also available for SCO XENIX and UNIX Systems!

It's multiuser and multitasking. And it's ready to give you true workgroup benefits while maintaining keystroke and file compatibility with Word for MS-DOS and OS/2, preserving your investments in Word training and data.

With Microsoft Word 5.0 for UNIX Systems, you can share a single copy of Word — on a single PC — with an entire workgroup of 16, 32, or even more users on inexpensive terminals.

And your workgroup can share documents, style sheets, forms, macros, glossaries, and outlines — plus group review and editing features such as annotations and redlining — while sharing expensive

Microsoft Word 5.0 joins SCO's powerful team of multiuser business applications.

printers and other resources as well.

Of course, Microsoft Word 5.0 for UNIX Systems still gives you Word's advanced productivity features. Data, text, and graphics integrated in a single document. Text flow around graphics. Side-by-side columns. A vast array of fonts and sizes. Links to spreadsheets and other programs. Line and box drawing. Spell checker and thesaurus. Speedkeys. And a whole lot more.

And since it's compatible with SCO Portfolio,™ Microsoft Word 5.0 for UNIX Systems supports the SCO Portfolio Clipboard, letting you copy and paste data to and from other popular applications such as SCO Professional and SCO FoxBASE+.

Call ext. 8605 at SCO today for the SCO Authorized Reseller nearest you. And see for yourself that if you want your SCO XENIX or UNIX System running the best multiuser word processing available today, you'll have our Word on it.

SCO®
THE SANTA CRUZ OPERATION

(800) SCO-UNIX (726-8649)
(408) 425-7222
FAX: (408) 458-4227
E-MAIL: ...!uunet!sco!info
info@sco.COM

SOFTWARE • PROGRAMMING

Generate Mapping Applications with TerraView

TerraView is a library of cartographic functions that you can use to integrate and display user data in map form for demographics, real estate, and other applications that manipulate spatial information. According to Terra-Logics, the library lets you treat maps and data as objects without concern for graphical screen operations or disparate map data formats.

TerraView can link symbols and text on a map to records in a database; applications can overlay, display, and operate on map data from the U.S. Geological Survey, Etak, Rand McNally, the U.S. Census Bureau, and other sources. TerraView functions, such as zoom, user-defined symbols, rotation, and scaling, are callable from many programming languages. Written in C, the library lets you write applications for DOS, Unix, the X Window System, Microsoft Windows, and VMS.
Price: $3495.
Contact: TerraLogics, 114 Daniel Webster Hwy. S, Suite 256, Nashua, NH 03060, (603) 889-1800.
Inquiry 1152.

An application created with TerraView responds to an inquiry about Central Park by highlighting downtown Manhattan.

Object-Oriented Pascal for the Mac

Symantec has released a new version of Think Pascal that supports Apple's MacApp library of classes for object-oriented programming in Pascal and includes Symantec's own Think Class Library, giving programmers a choice between a library that's comprehensive (MacApp) and one that's comprehensible (TCL).

A new class browser lets you examine and edit classes, while providing a way to navigate through complex object-oriented programs so that you can see how classes in a program are related. The browser shows a family tree for all the classes in a program so that you can see all the messages a class handles.

Think Pascal 3.0's linker supports segmentation directives that give you greater control over how programs are segmented. You can structure code in logical and functional units, and the linker strips out unused code down to the procedure level so that unused code portions don't add to the size of the final application. Symantec also improved the environment's resource editing capabilities, including resource tools licensed from Apple such as Rez, DeRez, and PostRez. Version 3.0 can now handle source files of up to 32,000 lines.

To run Think Pascal 3.0, you need a Mac Plus with 1 MB of RAM. You'll need 2 MB to run TCL and 4 MB to run MacApp.
Price: $249.
Contact: Symantec Corp., 10201 Torre Ave., Cupertino, CA 95014, (408) 253-9600.
Inquiry 1154.

A Simpler Look and Feel

A new version of Look & Feel, the designer utility that lets you prototype and simulate screens and turn them into C source code, is simplified and more portable than previous versions, Oakland Group says.

All the lower-level machine-dependent routines are isolated in the Oakland Windowing Library, so you can port the applications you build with Look & Feel 3.1 among different operating systems, such as DOS, OS/2, and Unix. Other improvements include fewer levels of indirection, smaller executables, and the ability to produce context-sensitive help. Look & Feel works with the company's C-scape Interface Management System, a programming library that helps you create data-entry, menu, help, and text screens, handling text and graphics together.
Price: $399 with C-scape; $149 without.
Contact: Oakland Group, Inc., 675 Massachusetts Ave., Cambridge, MA 02139, (800) 233-3733 or (617) 491-7311.
Inquiry 1156.

Identify Bottlenecks with Borland's Profiler

Borland says that the profiler recently added to its Turbo line is the first interactive profiler for PCs. With Turbo Profiler, you can highlight slow pieces of code and bottlenecks in programs, allowing you to rewrite slow portions or use different algorithms.

Turbo Profiler provides histograms, percentage values, and other statistics with which to measure the performance of selected parts of a program. You can examine code at the machine level, statement by statement, by groups of statements, by routines, or by defined areas. The profiler can map every time a selected piece of code is called, or it can take a statistical sampling. The menu-driven profiler can profile any program that has associated .MAP or CodeView tables.

Turbo Profiler 1.0, along with Turbo Assembler 2.0 and Turbo Debugger 2.0, constitute Borland's Turbo Debugger & Tools.
Price: Turbo Debugger & Tools, $149.95.
Contact: Borland International, 1800 Green Hills Rd., P.O. Box 660001, Scotts Valley, CA 95066, (800) 543-7543 or (408) 438-8400.
Inquiry 1153.

Graphics for Project Management Systems

Project management systems, while good for organizing schedules and tracking tasks, often provide only rudimentary graphical output at best. A new program called Graneda Personal uses a menu interface to work with project management systems such as Superproject Expert, Time Line, Artemis, Project Workbench, and Microsoft Project to create four types of drawings.

American Netronic says that a chart produced with Graneda Personal will have up to six times fewer crossovers than a drawing produced by a project management system.

The program runs on the IBM PC with 640K bytes of RAM and a hard disk drive.
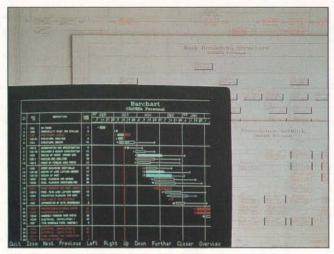**Price:** $995.
Contact: American Netronic, Inc., 110 Newport Center Dr., Newport Beach, CA 92660, (714) 760-8642.
Inquiry 1157.

Business Graphing on the Mac Without Typing Numbers

Tree Star's newest version of FlexiGraphs, the program that lets you create and edit business graphs without typing numbers into tables, includes a polynomial curve-fitting function. Once you draw a curve in a graph, FlexiGraphs 1.1 will generate a polynomial equation of that curve. The program also has new functions to smooth hand-drawn curves.

FlexiGraphs works in conjunction with Excel, WingZ,

The four types of drawings that Graneda can create include precedence network, arrow network, bar charts, and work breakdown structures.

Cricket Graph, and Delta-Graph by taking a different view of presentation graphics. Unlike other packages that require you to enter and change the data to create and modify charts, FlexiGraphs lets you create a spreadsheet from a graph. By changing the graph (i.e., drawing directly on it with a mouse), you change the numbers "behind" the graph that go into the spreadsheet, letting you perform what-if and break-even analyses intuitively.

FlexiGraphs 1.1 runs on the Mac Plus.
Price: $179.
Contact: Tree Star, Inc., 1802 Hillside Rd., Santa Barbara, CA 93101, (805) 682-4096.
Inquiry 1162.

Integrated Business Management on the Mac

Colleague 2 integrates scheduling, pricing, billing, job traffic, income and expense reports, and other business management functions. You can customize its screens and fields in the database. The program has a built-in mail-merge function.

Other features include graphing for reports, general ledger, sales reports, accounts receivable, and a daily calendar.

The program includes a run-time version of 4th Dimension and requires a Mac Plus with at least 2 MB of RAM.

Price: $395.
Contact: Colleague Business Software, Inc., 10000 Research Blvd., Suite 210, Austin, TX 78759, (512) 345-9964.
Inquiry 1161.

DRI Eases Presentation Pain

The new version of Digital Research's Presentation Team offers text charting, drawing, and graphics capabilities. Version 2.0 gives you desktop publishing–type control over word charts and a preview window that lets you see—all on one screen—the changes that take place in your image as you select different attributes.

You can create and edit presentations with all visuals in one file so that you need not close one file and reopen another to view the next slide in a presentation.

The program supports GEM, CGM, PIC, EPS, and several other file formats. The program runs on the IBM PC with 640K bytes of RAM.
Price: $495.
Contact: Digital Research, Inc., P.O. Box DRI, Monterey, CA 93942, (408) 649-3896.
Inquiry 1158.

New Spreadsheet Runs Under the X Window System

The eXclaim! program, which Quality Software says is the first spreadsheet to run under the X Window System, uses the OSF/Motif graphical user interface and takes advantage of X Window's mouse-driven windowed interface, not just the X terminal emulator. The Motif interface is carried throughout as eXclaim! uses its own three-dimensional pop-ups and icons. The new program has all the standard features of other spreadsheet programs like Lotus 1-2-3 and Microsoft Excel, including graphs and an alternate command interface bound to the Lotus-standard / (slash) key.

Fonts are adjustable, and Quality Software says that PostScript and LaserJet printers will produce near-WYSIWYG output when printer and screen fonts can be matched.

eXclaim! is for Unix System V 3.2 and 4.0 systems, Xenix, and OSF/1.
Price: $1000 to $10,000.
Contact: Quality Software Products, Inc., 5711 West Slauson Ave., Suite 240, Culver City, CA 90230, (213) 410-0303.
Inquiry 1159.

Finally. An input device based on your input.

Introducing SummaSketch® II.

New Lifetime Limited Warranty

The new SummaSketch II tablets were created with one thing in mind—you, the people who use tablets every day. You said you wanted a complete plug and play package, so we're giving you the works—both in PC and Macintosh® SE and II versions. A 12" × 12" or 18" × 12" graphics tablet with a 4-button cursor and 2-button stylus, or 16-button cursor for the PC.

The PC version includes interface cables for the IBM® PC, AT, PS/2 and compatibles. A utilities diskette with test and reset software, an Autodesk® Device Interface™ driver, Universal Mouse Emulator™ and a Microsoft® Windows driver. And an offer for

a free tablet template (US and Canada only) worth over $245.
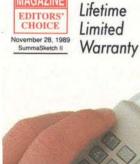
The Macintosh version has an Apple® Desktop Bus™ interface device to connect the tablet to the computer.

You'll also get the most software compatibility with over 250 PC programs and all Macintosh SE and II software written under the Apple Software Developers guidelines.

SummaSketch II tablets have a standard accuracy measurement of ±0.015 inches, selectable resolution of up to 1,016 lines per inch and high proximity so you can trace from documents up to ½" thick. Add in convenience features such as a power/proximity light, on-off switch, wedge shape design for easy use, lightweight construction for portability—and it's easy to see why SummaSketch is the industry standard and the

obvious choice of today's computer professionals.

Best of all, you get all of these benefits at an affordable price. And that's why our new SummaSketch II is the easiest buying decision you have to make. Find out more about SummaSketch II today. For literature and the name of a local dealer call 1-800-888-2028, Ext. 304. For technical information call 203-881-5400.

Summagraphics™

Every decision should be this easy.™

Finite Element Analysis on the Mac

LapCAD Engineering's new version of LapCADII, a finite element modeler that runs on the Mac, supports model sizes of up to 10,000 nodes. A PICT interface lets you run a blueprint through a scanner, read the resulting PICT file, reproduce the drawing, and mesh it into a finite element or solids model.

LapCADII 4.0 serves as a preprocessor and postprocessor for two finite-element-analysis programs, MSC/pal and MSC/pal 2, developed by MacNeal-Schwendler. Once the model is analyzed by the MSC programs, you use Lap-CADII to view the shapes and stresses.

LapCADII 4.0 supports wire-frame, hidden-line, shrunken-element, and shaded plots, and graphical modeling operations such as copy, rotate, snap, and contract.
Price: LapCADII 4.0, $495; MSC/pal, $1495; MSC/pal 2, $2495.
Contact: LapCAD Engineering, 1568 Diablo Point Court, Chula Vista, CA 92011, (619) 421-1236.
Inquiry 1163.

LapCADII, analyzing the stress contours of an automotive suspension part.

Real-Time Process Control

With Quick-Stat, a statistical control system, you can monitor manufacturing operations in real time, allowing you to eliminate waste and reduce downtime on an assembly line, Custom Technology says. The program is currently available for the IBM AT, and a Unix version with an X Window System display driver will be available by mid-June.

Quick-Stat can continuously acquire data from intelligent devices and immediately display data points against pa-

rameters such as an upper or lower control limit. In this way, the program can notify you when something on the assembly line is out of kilter.

As the program acquires the data from a programmable logic controller or any other measurement instrument, it can automatically calculate statistical functions. The program can simultaneously monitor up to 20 processes.

Quick-Stat supports devices from Alan Bradley, Gould-Modicon, Texas Instruments, other device manufacturers, and any device with an RS-232C serial port. It requires 640K bytes of RAM and a hard disk drive.
Price: Base configura-

tion, $1995.
Contact: Custom Technology, 300 West Grand, Suite 605, Chicago, IL 60610, (312) 384-4104.
Inquiry 1164.

Emulate Pen Plotters with PrintAPlot

A new version of Print-APlot, the utility that lets your printer emulate a conventional pen plotter, can print large drawings over multiple pages and lets you position an image anywhere on a page.

Other features of version 2.0 include rotation of drawings by 90, 180, or 270 degrees; the ability to scale drawings from 1 percent to 600 percent of original size; independent *x* and *y* positioning; and a generic HPGL setting that lets the printer emulate more plotters.
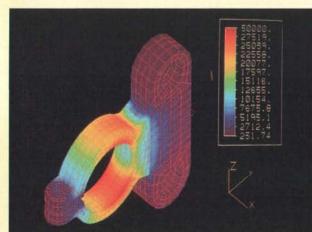
PrintAPlot 2.0 consumes 234K bytes of RAM.
Price: $299.
Contact: Insight Development Corp., 2200 Powell St., Suite 500, Emeryville, CA 94608, (800) 825-4115 or (415) 652-4115.
Inquiry 1167.

A Kinder, More Accurate, Nonlinear FEA

AccuPak, a new package for Algor's Design and Finite Element Analysis System for the IBM AT, combines a decoder and precision contouring with stress-to-nodes capability to make it easy for design engineers to perform accurate analyses. AccuPak's decoder uses a spreadsheet interface instead of providing you with lengthy text files. With the decoder, you can specify information about the model that you aren't able to specify in the modeling phase.

The stress-to-nodes technique extrapolates nonlinear

stress values to element nodes and edges, where maximum stress is most likely to occur. Independent estimates of stress values at shared nodes are compared to indicate accuracy; from the comparison, the program assigns a confidence factor that you can plot to compare with stress plots.
Price: AccuPak, $895; base system with ViziCad Plus, $889.
Contact: Algor Interactive Systems, Inc., 260 Alpha Dr., Pittsburgh, PA 15238, (412) 967-2700.
Inquiry 1165.

An Avalanche of CD-ROM Information

We've recently been buried by an avalanche of releases describing CD-ROM subscriptions and HyperCard directories of services. Here are a few of them.

CD Guide, the magazine that covers the audio compact disk field, is now offering a quarterly CD-ROM subscription that provides over 40,000 listings, 4000 full reviews, sample audio tracks, and liner notes of CDs. Each update features new reviews, listings, and about 10 disks that include a sample track.

In addition to serving CD cognoscenti, the CD-ROM—also available for touch-screen monitors—can act as a point-of-purchase tool.
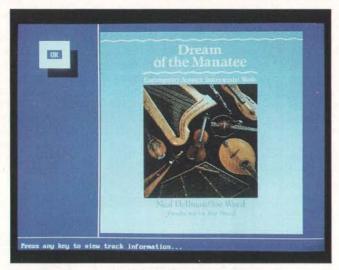
Filters let you search for a CD by several categories, including name, title, disk label number, and, if the disk is reviewed, sound and performance ratings. The program also lets you play any CD on your CD-ROM drive.

CD Guide Optical Edition conforms to the ISO 9660 standard and runs on an IBM PC.
Price: Single issue, $99; quarterly subscription, $259.
Contact: WGE Publishing, Inc., Forest Rd., Hancock, NH 03449, (800) 227-1053 or (603) 525-4201.
Inquiry 1168.

Shakespeare on Disc! includes the bard's 37 plays, five narrative poems, and 154 sonnets in both the Queen's English and American English versions. The CD-ROM requires an IBM XT with 640K bytes of RAM.
Price: $99.
Contact: CMC Research, Inc., 7150 Southwest Hampton, Suite C-120, Portland, OR 97223, (503) 639-3395.
Inquiry 1169.

With the CD Guide Optical Edition, you can examine liner notes or cover art while listening to the best cut from a featured disk.

MicroPatent is making 17 years' worth of abstracted information on U.S. patents available to IBM PC users who subscribe to its Automated Patent Searching (APS) system.

APS runs on the IBM XT or higher with 640K bytes of RAM, a hard disk drive, and a CD-ROM drive.
Price: Weekly subscription, $3150 per year; monthly, $950 per year; two-year back-file disks, $90 each; Toshiba XM 3201 CD-ROM drive, $895; Hitachi CDR 3600 drive, $799.
Contact: MicroPatent, 25 Science Park, New Haven, CT 06511, (800) 648-6787 or (203) 786-5500.
Inquiry 1170.

AHyperCard-based directory called the Macintosh Yellow Pages provides a listing of Macintosh software, hardware, accessories, books, CD-ROM titles, and magazines. The program has an indexing system that lets you search for a product using synonyms.

Each product listing includes a one-sentence description of the product. As in the Yellow Pages of your phone book, display ads purchased by companies further describe their products.

The Macintosh Yellow Pages will be updated on a regular basis. The program requires HyperCard 1.2 and a Mac Plus.
Price: $39.95.
Contact: Digital Media Publications, Inc., 362 Memorial Dr., Suite 121, Cambridge, MA 02139, (617) 225-7121.
Inquiry 1171.

The Specialty Series of databases from Aries Systems contains journal references selected from the National Library of Medicine's Medline database with abstracts drawn from more than 3000 biomedical journals. Initial offerings include clinical and research pathology, radiology, and primary clinical care.

Disks are available in High Sierra and ISO 9660 formats and run on the Mac Plus or higher (a Mac SE or higher is recommended).
Price: PathLine (pathology) and RadLine (radiology): $995 for a quarterly subscription, $595 for an annual subscription. Core Journals (clinical and primary medicine): $795 for a quarterly subscription.
Contact: Aries Systems Corp., One Dundee Park, Andover, MA 01810, (508) 475-7200.
Inquiry 1172.

Programmable DOS Shell for Power Users

Mi-Shell is a cross between a programmable Unix shell and a DOS shell such as Norton Commander, OpeNetwork says. The shell lets you browse through directories and execute commands on files or sets of files with a point-and-shoot or command-line interface.

The shell's programmable functions let you streamline your routines. You can assign arbitrary sequences of actions, written as a Mi-Shell script, to any keystrokes. Actions can depend on the environment, such as the file you are currently pointing to or its extension.

Mi-Shell lets you modify the size, appearance, and content of the display panels. You can display more or less information on each file and sort files in many ways, including by size, extension, and time stamp.

A built-in debugger is provided for the Mi-Shell script, which is a Forth-like language. Mi-Shell requires 130K bytes to initialize and needs 10K bytes of RAM when it executes an application.
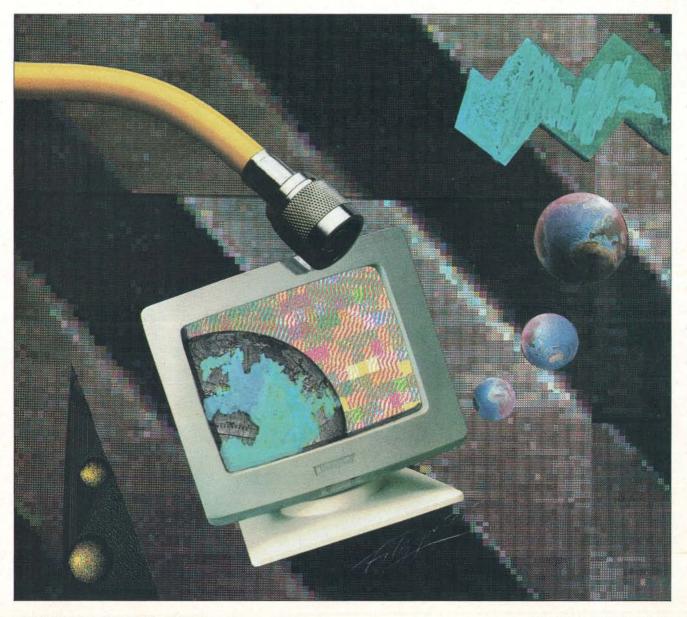
Mi-Shell integrates with two other power tools from OpeNetwork: the Berkeley Utilities, which lets you execute Unix commands from your DOS environment, and Delta, an advanced file-comparison utility that is suited for programmers. You can use Mi-Shell instead of the command line to invoke these tools.
Price: $89; Berkeley Utilities, $125; Delta, $79.
Contact: OpeNetwork, 215 Berkeley Place, Brooklyn, NY 11217, (718) 638-2240.
Inquiry 1174.

BYTE

INTERNATIONAL

Video Analysis on a PC

John Unger

Java brings minicomputer-style image processing to the desktop

I f you're involved in image analysis and you are on a budget, you ought to consider Java. It's a comprehensive software package for image processing and analysis, and it's one of the first to run on a PC.

Image-analysis programs perform an almost endless variety of analytical operations on raster graphics images. Raster images typically include those produced by scanners, paint programs, video cameras, or instruments that convert laboratory measurements into images (CAT scans and x-rays fall into this category). You can use Java, for example, to automatically delineate and measure the shoreline of a lake on a map and then calculate the lake's area. Or you could analyze an aerial photograph and count all the houses in the image. Or a medical researcher could use Java to count cells or compute bone density by measuring image intensities. (In fact, all the sample images included with Java are from the medical field.)

Traditionally, scientific image analysis/processing has been the domain of the large, powerful minicomputers and workstations. Although the PC-based Java does not seriously threaten these platforms, it provides almost all the larger machines' functionality and analysis capabilities at a far smaller investment in hardware and software.

Desktop Requirements

Java requires an 8086-, 286-, or 386-based MS-DOS computer with 640K bytes of RAM, a graphics card, a monitor, and a mouse or other pointing device. It also requires a frame grabber (a video digitizing board) to hold the images and an RGB or composite monitor (in addition to the computer monitor) to display images. Finally, Java requires a video camera or other device to capture the images you want to analyze.
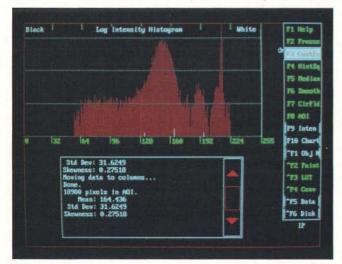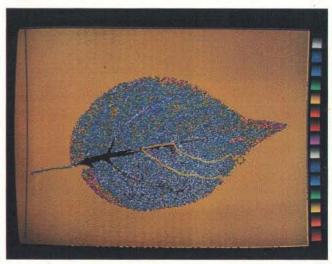
Java supports 640- by 200-pixel-resolution two-color CGA, EGA, VGA, Hercules monochrome and InColor, and AT&T 6300 400-line modes. Java takes advantage of an 80x87 math coprocessor, but it does not require one.

I tested Java on a Dell System 200 12.5-MHz 286 system with VGA color, a Mouse Systems optical mouse, and a PCVisionplus Frame Grabber board connected to an Electrohome ECM 1312U high-resolution color monitor and Cohu solid-state video camera.

A fully configured image-processing system is not cheap. The Java software alone sells for $1495. The PCVisionplus Frame Grabber costs $1995, the RGB monitor retails for $1295, and the high-resolution video camera and lens cost $1770. That's a total cost of $6555, not counting your PC. These costs are still much less than a minicomputer-based image-processing system.
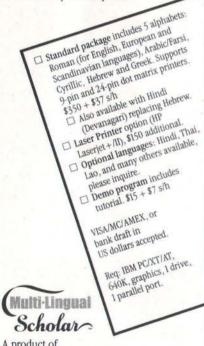
Java uses the computer system's monitor to show program prompts, and a second image monitor for the processed image.

Java 1.30

Company
Jandel Scientific
65 Koch Rd.
Corte Madera, CA 94925
(415) 924-8640

Hardware Needed
IBM XT, AT, PS/2, or compatible with 640K bytes of RAM and CGA, EGA, VGA, or Hercules graphics; video frame grabber board; mouse or other pointing device; separate RGB or composite monitor for image display; optional digital video camera for image capture

Price
$1495

Inquiry 886.

Jandel recommends that you purchase Java and the required hardware components from one of its local dealers, who can install the system in your computer. This is a good idea. As I discovered, getting a frame grabber board properly installed and configured in your computer can be difficult. Fortunately, Jandel has on-line technical help available via a toll-free number, and its engineers are familiar with potential pitfalls and problems. They solved my hardware installation troubles skillfully, and I didn't have to wait for someone to call me back. The dial-up help is available to all registered Java owners.

Java uses about 1.5 megabytes of space on your hard disk drive, less if you don't load the images used in the tutorial. The Java manual is divided into two main sections. The first describes some of the general concepts of image processing and analysis, tells how to install and set up the software, and takes you through six step-by-step tutorial sessions using images supplied with the software. The second section is a reference to all of Java's functions and commands and numerous useful examples.

On-line help is available for any menu selection or at any prompt or request for data in the dialogue window. Java also has an organized menu structure that minimizes the amount of moving around when working on a specific stage of processing or analysis. Also, Java can memorize the steps of your analysis in a macro file, which you can play back to repeat an often-used process.

Dual Images
To use Java, you must work with both the image-display monitor and your com-puter's system monitor. Java shows menu choices and program responses on the system monitor. You can use the mouse for both making menu choices that appear on the computer monitor and controlling the cursor that appears on the image-display monitor. You can set the mouse sensitivity by editing a start-up configuration file. However, I found it difficult to find a setting that was responsive when moving around the menus on the computer screen, yet was not too sensitive on the image screen.

Java's main program screen (on the system monitor) consists of a vertical menu of the program's major functions, a scrollable dialog box, and an area for graphics, such as histograms.

Java lets you issue file-related operating-system commands to retrieve and store files, change directories, and so forth. It also provides a data worksheet of 16,384 columns by 65,372 rows that you can use to store, display, reorganize, and analyze the measurements and data from an image. You can save the data in the worksheet as an ASCII file and then import it into your favorite spreadsheet for further analysis. Interestingly, this built-in worksheet function is not available on any of the minicomputer-based image-processing software I have used; this feature added a new dimension to my ability to analyze image data.

There are two major ways of making images to be analyzed: (1) You can capture a photograph, map, graph, or other image with the digital video camera connected to the frame grabber board, or (2) you can use an image captured by someone else and saved as a disk file. Jandel has chosen to have Java save and store images only as 8-bit TIFF files. If your images are not in this format, you must convert them to an 8-bit .TIF format before Java can use them. Most scientists use digitized images made by a wide variety of instruments and from the output of other computer programs. Java would be more accessible if Jandel furnished software to convert these other image formats to the .TIF format.

Java works with monochrome, gray-scale 8-bit images. In the case of the PCVisionplus Frame Grabber, these images are 640 by 480 pixels in size. Because each pixel is represented by 8 bits, it can have one of 256 unique values. Within Java, each pixel is mapped to a gray scale that ranges from 0 (black) to 255 (white) for display on the image monitor. This scaling will display the images generated by the video camera on the image monitor as if they were high-

quality black-and-white photographs.

I used a camera and lens that could capture an original image as small as about 6½ by 4¾ inches. At 640 by 480 pixels, that's a resolution of only about 100 pixels per inch. In other words, on that size image, the smallest entity the software could possibly recognize would be approximately 0.01 inch—or about three times larger than the dot made by a typical laser printer.

I used the video camera to capture images of colored maps—both topographical maps with subdued shades and geological maps with many different colors, as well as monochrome images and charts. Not surprisingly, the gray shades on the captured image that represented some of the colors on the original maps were difficult to separate, because

J ava

provides a tremendous amount of analyzing power for the money.

they lacked contrast. Even using Java's processing tools, it was hard to discriminate between the areas that had different colors on the original. I could, however, examine the intensities of those areas. These measurements showed that very little and, in some cases, no discernible contrast existed between some of the original image's colored areas.

In short, Java analyzes pixel intensities of an image. That data may not directly correspond to what you see. Frequently, the intensities in an image are proportional to some measurement, such as the amount of scattered sonic energy (i.e., a sonogram) or the total amount of transmitted gamma radiation (i.e., an x-ray). For example, Java can analyze a digital elevation map of a mountain. The intensity of each pixel in the image would be proportional to the elevation of the ground surface scaled to a value between 0 and 255. Java can map these pixel values to a gray palette that depicts the pixels with the largest values as the lightest shades and the least values as the darkest. Then the palette depicts the highest elevations as white areas on the image monitor, and the lowest as black.

Java uses color, though, to let you map the intensity values of the pixels in an image to an arbitrary color palette instead of to a gray-scale palette. This operation results in a pseudocolor rendition of the image, in which a specific color is associated with a given pixel intensity. Using the example of the elevation map of the mountain, you could use a palette where 256 colors ranging from violet to red (the inverse of the spectrum) are assigned to the pixel intensities or elevations. Now the picture of the mountain would show the lowest elevations as cold or blue colors and the highest elevations as warm reds and oranges.
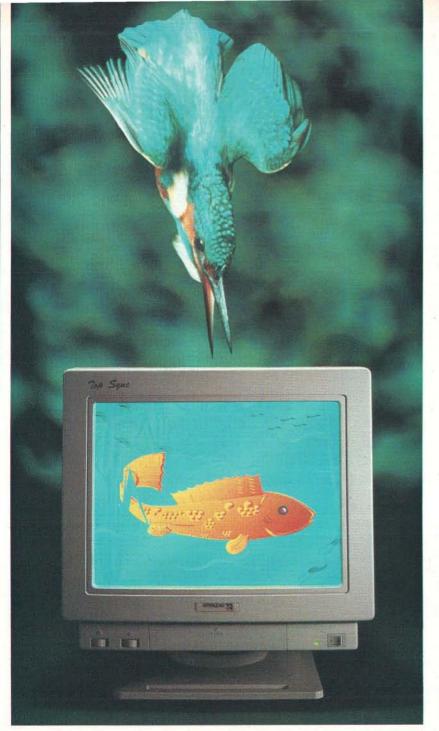
Real Analysis

Java provides honest-to-goodness image analysis and processing on a PC-based microcomputer. Its only real shortcomings are a result of the small range of intensities it can represent in an image (8 bits per pixel), and its maximum image size of 640 by 480 pixels. Java also cannot handle true RGB color images with 3-bit planes of data or 24-bit intensity data, the norm for larger systems running on minicomputers. On the other hand, Java provides a tremendous amount of analyzing power for the money, and its processing and analysis functions, while not as complex, rival those on larger systems.

I found Java's speed acceptable while performing its more complex processing and analysis functions on my 286 machine. For example, running a high-pass filter to enhance the edges in an image took 15 to 20 seconds. That didn't seem much slower than the larger systems I am familiar with, considering that Java is analyzing a smaller-size 8-bit image compared to a larger 24-bit image for the minicomputers. I was bothered, though, by the software's somewhat sluggish mouse response while making measurements on the image display.

A fast 386 system combined with Java would make an excellent image-analysis system for a small research lab where the main tasks entail working with small-size monochrome images. Many scientists are already using microcomputers for other research chores, so it is only a matter of adding the frame grabber, the video camera (if necessary), and the Java software. Otherwise, it would be more difficult for someone who is used to working with one of the larger, 24-bit color image-processing systems to accommodate his or her research to Java's more limited capabilities. ∎

John Unger is a geologist for the U.S. government in the Washington, D.C., area. He can be reached on BIX as "junger."

OUR BRILLIANT DISPLAYS CARRY A CLEAR MESSAGE.

Our range of data monitors is a brilliant collection of highly versatile units - each one ready to give you a precise and vivid display.

InterQuad provide exactly what you are looking for; sharp images, crisp text, and vivid colours, all presenting the information you need with faultless clarity.

Here is the reliability and the performance you require to meet present and future information display needs to perfection.

'Quality' is the ultimate message which shines out clearly from every screen. These display monitors are designed and manufactured to the highest standards, with every unit passing a series of stringent test procedures.

They also have 'the right attitude' to everyday working - all feature an integral tilt and swivel base, permitting choice of the best viewing position for comfort and convenience. The result: monitors which will be an asset to any office.

<u>A BRAND NEW 15"!</u> *Look closely at our new autoscanning / TOP SYNC, flat screen 15- inch intelligent high resolution monitor! (Main illustration). It automatically synchronises 30 to 50 kHz and 50 to 80 Hz, compensating for picture size variation. It's compatible with all IBM PC and PS/2 graphics, including VGA, and supports 1024 x 768 non-interlaced.*

InterQuad can give you a completely fresh angle on brilliant graphics, data, and whatever you want to see.

Contact us at the addresses below.

CM 1401
14" EGA with 640 x 350 resolution in EGA mode and 640 x 200 in CGA for XT/AT.

MS 1422
14" Multiscan colour monitor for XT/AT and PS/2 families of Personal Computers.

AM 1412
14" Analog Mono high resolution 720 x 480 ideal for DTP and WP type applications.

AC 1432
14" Analog Colour VGA monitor compatible with all VGA systems including the IBM PS/2 range.

INTERQUAD LIMITED

GRASP THE FUTURE

The road ahead is 32-bits wide, and Twinhead has the map

Superlap-386SX

A portable handful of power and performance. Features 386SX, RAM up to 4MB, FDD and HDD, VGA graphics with 32 greys, 640×480 page-white LCD, 83-key detachable keyboard with embedded keypad. Ports include 2 serial, 1 parallel, external monitor/disk drive/keyboard/keypad, expansion bus, and interface for an internal modem. Using AC power or NiCad battery, Superlap gives you 32-bit power on the move.

Superset-600

A compact platform for fast 386 or 486 processing. The basic 600 (386-25MHz) is slotted for simple Power Card upgrades. Power Cards supply CPU (386-25/33, 486-25/33), 64KB cache, and cache controller. All-in-One design integrates VGA, parallel and 2 serial ports, and drive controllers on mainboard. Storage for 2 × FDDs, 3.5″ HDD, and 5.25″ Tape/HDD. Superset-600 is your bridge to the future.

Staying abreast of everchanging technology is critical. Twinhead has an unparalleled record of engineering the latest ideas into reliable, affordable systems. Take the future in hand, see Twinhead's new 32-bit models. Your systems for the nineties.

Twinhead

Engineering Ingenuity

Global Industrial Center
2nd Floor, 2 Lane 235,
Bao Chiao Road,
Hsin Tien, Taiwan, R.O.C.
TEL: (886-2) 917-9036 Fax. (886-2) 917-2674/5

Opus Datasafe 120

Dick Pountain

The Opus Datasafe 120 is one of the first, if not the first, PCs to be designed with data-security mechanisms built into both its hardware and its system software. It is an AT compatible powered by a 12-MHz 286 CPU and fitted with twin 60-megabyte hard disk drives that can be configured as either 120 MB on two hard disk drives or a single mirrored 60-MB hard disk drive.

A few years ago, many people would have regarded data security as something that only the mainframe users worried about, perhaps even as paranoiac. But as I'm writing this review, the ripples of the "AIDS Information Disk" scandal are only just subsiding.

You might remember that incident: 10,000 copies of a phony disk containing a Trojan Horse program were mailed to banks, hospitals, and research laboratories in the U.K. and Europe, causing some users to lose data on their hard disks and creating a great deal of media hysteria. That deleterious disk could not have been loaded onto an Opus Datasafe, except by a supervisor. Perhaps one day we will look back fondly to the golden days when people could swap disks freely, before the Serpent gave Adam and Eve a BIOS listing and a map of the boot sector.

Of course, viruses and Trojan Horse programs are not the only, or even the most important, threat to data on PCs. More people lose data every day from hard disk malfunctions than they ever do from a virus. Opus has covered this angle, too, by providing shadowing (sometimes called mirroring) for its hard

A system with built-in data-security mechanisms that has advantages and disadvantages

disk system. Data is written to two identical but independent hard disk drives, which are very unlikely to fail simultaneously. Clever BIOS software makes this process totally transparent to DOS and leaves the machine fully DOS compatible.

Opus is a U.K. firm that has made a name over the last few years for offering well-made and keenly priced PC clones. Opus computers are wholly designed and built in the U.K.—an unusual situation, nowadays. The Datasafe is Opus's most adventurous hardware design so far.

Hardware Overview
The Datasafe is by no means a compact machine; its case is roughly 17 inches square and 6 inches deep. It is, however, attractively designed and of a much higher quality than a run-of-the-mill clone. The case is made of plastic-coated steel with a stylishly molded front panel that bears a security lock, disk and power lights, and buttons marked Reset and Turbo. The Turbo button is redundant, because the Datasafe always runs at 12 MHz.

Fitted to the review machine are a high-density 5¼-inch and a 3½-inch floppy disk drive. It also has a 102-key Enhanced keyboard with a good clicky feel, although its weight is too light for my taste.

The advantage of the large case is that it can accommodate six AT-bus slots (four 8-bit and two 16-bit). A feature-packed motherboard means that the system uses only one of these, leaving five free for the user's cards.

Mounted on the rear panel
continued

Company
Opus Technology Ltd.
53 Ormside Way
Holmthorp Industrial Estate
Redhill
Surrey RH1 2LW
UK
44-737-765080

Components
Processor: 12-MHz 286
Memory: 1 MB of zero-wait-state RAM
Mass storage: Twin 60-MB hard disk
drives; machine can be configured with
5¼-inch and 3½-inch floppy disk drives
Display: Philips color VGA monitor
Keyboard: 102-key Enhanced layout
I/O interfaces: Two RS-232C serial
ports (9-pin and 25-pin); one bidirectional
parallel printer port; IBM Enhanced
keyboard–compatible keyboard socket;
analog VGA video output

Size
(System unit only): 17 × 16½ × 6⅖
inches; 20 kilograms

Price
With monochrome VGA: £1999
With color VGA: £2499

Inquiry 882.

are the main power connector (with a chained outlet to drive a monitor), two serial ports, a parallel printer port, and the keyboard socket; a large and rather loud fan cools the interior. Opening the case provides the first evidence of Opus's security; you must turn the key in an AT-style case lock and then use the special screwdriver supplied to remove the five retaining screws, which an ordinary crosshead will not shift. Stealing cards from this machine is a job for a professional.

One card that would be worth stealing is the VGA card, which occupies one of the 8-bit slots. Rather than implementing a VGA on the motherboard with one of the several available chip sets, Opus has chosen an add-in unit from Tseng Labs called the EVA/800. It is a 256K-byte, 8-bit VGA card with extra features like hardware-assisted windowing, zooming, and panning.

Ordinary PC software packages cannot exploit these nonstandard features unaided, but Tseng Labs provides special drivers for the popular programs AutoCAD, Lotus 1-2-3, Windows/286, Ventura Publisher, and WordPerfect 5.0

on the utility disks. The monitor supplied with the system is a Philips color VGA monitor badged with Opus's name, which gives a display with a good color and sharpness.

Also included is a memory-resident, hot-key-activated Zoom and Pan utility that operates from the numeric keypad. Opus claims that this utility works with any graphics package that uses BIOS interrupts; the only program I could make it work with was PC Navigator, and it hung the machine completely with Flight Simulator. In any event, zooming in on a bit-mapped screen produces an image resembling a crossword puzzle and is of little practical use.

Perhaps of more interest to the average user is that I found the EVA/800 to be one of the most compatible VGA cards I have yet seen. For example, it ran all the tweaked VGA modes up to 800 by 600 pixels by 16 colors in the public domain fractal program FRACTINT (on my multisync monitor). This suggests that it is fully register-compatible with IBM's VGA.

A neat trick of the EVA/800 is that it can convert CGA, Hercules, and Monochrome Display Adapter output to VGA frequencies so you can display software using these older standards on a fixed-frequency VGA monitor instead of on an expensive multisync monitor.

Opus designed the Datasafe's motherboard around VLSI Technology's AT chip set; the board includes the hard and floppy disk drive controllers and the serial and parallel port hardware (the hard disk drive controller is rather special, but I'll discuss this later). The motherboard has a 12-MHz 286 CPU, and 1 MB of zero-wait-state RAM is mounted on the motherboard as standard. The BIOS uses the area above 640K bytes as extended memory for disk caching. You can expand RAM to up to 16 MB by adding an expansion card.

The ROM BIOS is based on American Megatrends' AMI 286 code, but Opus has modified it to support the Datasafe's hardware security lock and the shadowed hard disk. For example, the low-level format routines need to be able to format both disks together. The BIOS also supports ordinary AT hard disk drives, although these need a conventional controller in an expansion slot.

The usual power-on self test and CMOS configuration routines have been modified, too. On the Datasafe, the POST not only tests RAM and peripherals but also checks whether or not the contents of the two shadow disks are identical. If they aren't, the POST per-

forms a repair to make them so.

When you save files to the hard disk in the normal course of your work, they are written to both disks simultaneously (invisibly to you). Hence, the two disks should always contain exactly the same data. When you shut down the machine, which must always be done using Opus's Shutdown utility, a special record is written that says the disks' contents are identical.

If, however, the machine shuts down prematurely due to mechanical or electrical failure or a power cut, then this record doesn't get updated. The next time you boot the machine, the BIOS recognizes that an abnormal termination has occurred and tries to repair any data loss.

Of course, you can also provoke this repair unnecessarily by forgetting to use the Shutdown utility and just switching the system off, as I did several times. The repair operation can take 10 minutes or more on a full disk, but you can abort it by pressing Ctrl-Alt-Del in case of a false alarm. That same action has been modified so that it offers a menu of choices: shutdown, reboot, or cancel.
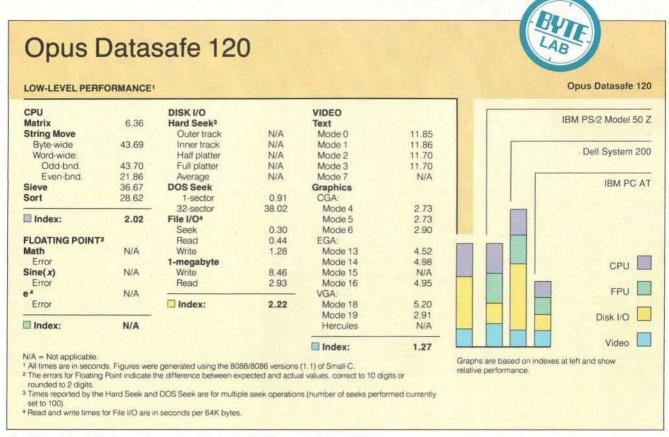
You can invoke the CMOS configuration utility by pressing the Delete key during boot-up. It contains the usual routines to set the battery-backed clock/calendar and the hard disk type and to do a low-level format.

This utility also contains Opus's own routines to copy from the active hard disk to its shadow twin and vice versa and to relocate bad sectors on the disks. If one of the hard disks were to suffer a head crash, you would replace it and then use this same utility to restore the data from its shadow. There is also a routine that lets you configure the shadow disk pair as two separate disks, should you need the extra space.

Note that disk shadowing cannot protect against operator errors (like an unintentional DEL *.* or FORMAT C:), either of which will destroy the data on both disks simultaneously, but you can prevent this disaster with software—an issue I address in the next section.

Security Features
The Datasafe has a hardware security lock quite separate from the case lock on the front panel (which most AT clones have nowadays). It is mounted at the back of the case and uses a high-security key that is not easy to copy. This lock is connected to the electronics so that the BIOS can detect the lock's state. Just for good measure, it acts as a second physical case lock, too, backing up the case lock on the front.

Opus Datasafe 120

LOW-LEVEL PERFORMANCE[1]

Opus Datasafe 120

CPU		DISK I/O		VIDEO	
Matrix	6.36	**Hard Seek[3]**		**Text**	
String Move		Outer track	N/A	Mode 0	11.85
Byte-wide	43.69	Inner track	N/A	Mode 1	11.86
Word-wide:		Half platter	N/A	Mode 2	11.70
Odd-bnd.	43.70	Full platter	N/A	Mode 3	11.70
Even-bnd.	21.86	Average	N/A	Mode 7	N/A
Sieve	36.67	**DOS Seek**		**Graphics**	
Sort	28.62	1-sector	0.91	CGA:	
		32-sector	38.02	Mode 4	2.73
▢ **Index:**	**2.02**	**File I/O[4]**		Mode 5	2.73
		Seek	0.30	Mode 6	2.90
FLOATING POINT[2]		Read	0.44	EGA:	
Math	N/A	Write	1.28	Mode 13	4.52
Error		**1-megabyte**		Mode 14	4.98
Sine(x)	N/A	Write	8.46	Mode 15	N/A
Error		Read	2.93	Mode 16	4.95
e^x	N/A			VGA:	
Error		▢ **Index:**	**2.22**	Mode 18	5.20
				Mode 19	2.91
▢ **Index:**	**N/A**			Hercules	N/A
				▢ **Index:**	**1.27**

N/A = Not applicable.
[1] All times are in seconds. Figures were generated using the 8088/8086 versions (1.1) of Small-C.
[2] The errors for Floating Point indicate the difference between expected and actual values, correct to 10 digits or rounded to 2 digits.
[3] Times reported by the Hard Seek and DOS Seek are for multiple seek operations (number of seeks performed currently set to 100).
[4] Read and write times for File I/O are in seconds per 64K bytes.

IBM PS/2 Model 50 Z
Dell System 200
IBM PC AT

CPU
FPU
Disk I/O
Video

Graphs are based on indexes at left and show relative performance.

For a full description of all the benchmarks, see "Introducing the New BYTE Benchmarks," June 1988 BYTE.

When the hardware security lock is on, three things happen. First, the BIOS intercepts all Control-Break and Control-C key presses. Because of this feature, you can't break out of the AUTO-EXEC.BAT file at boot time and gain access to the DOS command line. Second, the Datasafe can only be booted from the hard disk; at boot time, it ignores any floppy disks. This process blocks the other route you might use to get access to plain, unprotected DOS. Finally, access to the BIOS utility routines becomes password protected.

With the lock on, the Datasafe executes the contents of AUTOEXEC.BAT on the hard disk. Only an electronic engineer with a good hacksaw can prevent this process from taking place. If you have set up Datasafe for full security, then AUTOEXEC.BAT loads a software security shell, which provides the next level of security.

The shell, which Opus supplies as standard, is called Certus Junior from the U.S. firm Foundation Ware; the Junior means it is a cut-down version of a larger system. Under Certus Junior, you must have your own password to use Datasafe, and, having logged on, you are allowed access to only those directories and executable programs that the system

supervisor has chosen for you.

With Certus installed, you boot up the Datasafe, and a log-in window appears, into which you must type a valid user name and password. You then get access to that user's programs and directories, but, at the same time, a history log file records your user name, whether or not your log-in succeeded, and how many attempts you made.

If, after three tries, you cannot log in correctly, the Datasafe locks up, and you must reboot it. Any attempts to fiddle about with other people's accounts are permanently recorded for the supervisor to see. In addition, Certus keeps a daily log that shows which programs each user ran and for how long.
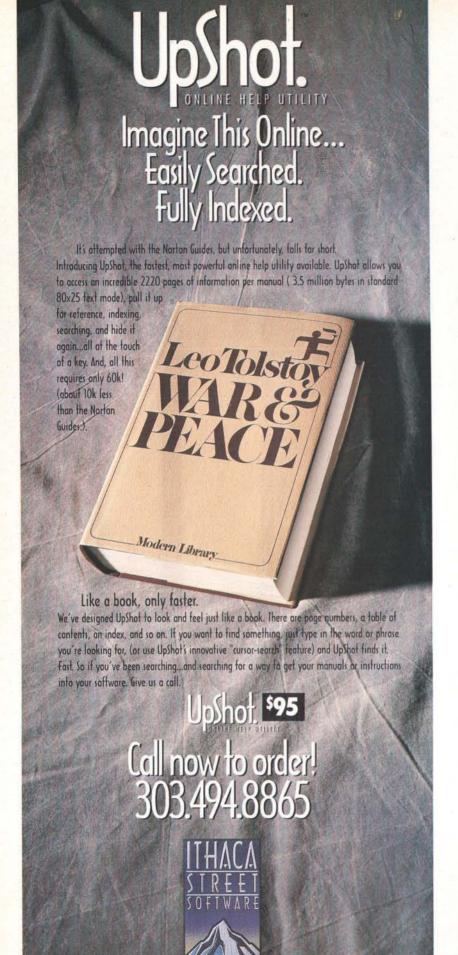
With Certus running, it is impossible for anyone other than the supervisor who has the security key to load any new software into the Datasafe. To add new programs, you need to turn off the lock, reboot from a floppy disk, copy the software, and then register the programs with Certus for all those accounts that need to use them.

Certus is intended for use on shared machines in big offices where a few core applications are run all the time by different people. It prevents the bored office boy from running virus-ridden pirated

games just as it prevents someone from performing industrial espionage. It is very inconvenient to run Certus on a single-user machine, especially for someone like myself who needs to constantly add and remove software from the system. Recognizing this fact, Opus supplies the Datasafe without Certus installed, which encouraged me to test the installation procedure.

Installation

Certus is largely self-installing. You run the Install program from drive A, which copies the software onto the hard disk and modifies AUTOEXEC.BAT to run Certus on booting. The Install program also scans your whole hard disk to identify all the executable programs (and any overlays) and builds a sorted database of them. You must then log in as user "super" and run a program called Manager to create as many new user accounts as you need, with names and passwords. For each user, you are presented with an X-tree-style directory picture, and you can set the directories and files that each user may access interactively by pointing with the cursor.

The problem with any impregnable defense mechanism is that you might lock
continued

yourself out. During my first attempts to install Certus, it kept denying me access to everything, and I could only regain control by booting from a floppy disk. The problem was eventually traced to a shareware program I had loaded, which had the file date 3-10-2011. A date in the next century upset Certus's file scan so much that it locked up the whole system, but changing the date with Norton's File Date fixed the problem. This date sensitivity seems to be a reproducible bug in the current version of Certus.

Another problem I had with Certus was that I couldn't get SideKick and other pop-up programs to load properly from AUTOEXEC.BAT. You have to load pop-up programs before the Certus resident modules DIRLOCK, RESIDENT, and LOGIN, which must be the last things in AUTOEXEC. When it's loaded, SideKick pops up, but Certus intercepts all its keystrokes so that you cannot actually operate it at all. Another program, Tornado Notes, won't even pop up. Both work fine if you load them manually after boot-up (from a directory that you have permission for, of course), so it is purely an inconvenience rather than a disaster.

Of course, apart from Certus, you can use other shells for security. On a networked operating system like NetWare, you would use the operating system's own log-in and security features, with the Opus hardware ensuring that users cannot break out into plain DOS.

Performance
The 12-MHz 286 is now the standard engine for entry-level AT clones, so the system's computational performance must be considered average. The BYTE low-level benchmarks show it performing as expected: roughly twice as fast as an IBM AT and about 50 percent as fast as a 25-MHz 386 machine on computationally intensive tasks.

Disk performance is another matter. Due to a happy combination of design features, the Datasafe offers hard disk performance well beyond that of most similarly priced machines from other manufacturers. Rather than imposing overheads on disk performance, a well-designed shadow disk scheme permits an improvement in disk seek time, because the head of the second drive can be moved into position for the next read while the head of the first drive is reading a track. Since the data is identical on both hard disk drives, it doesn't matter from which one any given data record is read.

continued

"I Get Paid to Write Programs. And Now, I'm Writing Them Twice As Fast."

If you write programs for a living, get to know The Clarion Professional Developer® 2.1.

It can cut applications development time by 50% or more without compromising your design. It's so fast and easy, you can complete prototypes while your clients watch.

Professional Developer 2.1 can interface to routines from C and Assembler. It allows import and export of dBASE, BASIC and DIF files. Its application generator produces commented source code which can be easily adapted with its utilities to allow you to generate .EXEs. Network support is included at no extra cost. There are no run-time licenses or fees. And now Version 2.1 includes Report Writer, which allows you to easily and intuitively create ad hoc reports and queries (from Clarion files, non-Clarion files or any combination).

Professional Developer is an immensely powerful, easy-to-use DBMS application developer. We didn't say it: The PC WEEK poll of corporate satisfaction said it. We came out head and shoulders above the 15 other programmable databases. No one else came even close.

In this business where time is money, let Professional Developer help you generate more money for your time. Below is a partial list of our world-wide network of dealers. For the one nearest you, call:
(U.S.A.) 800-354-5444   (Canada) 305-785-4555

Here's what American corporations in a PC Week survey thought of the power and ease of use of Professional Developer vs. competitors. Any questions?

CLARIONSOFTWARE™

150 East Sample Road, Pompano Beach, FL 33064

Circle 413 on Reader Service Card (DEALERS: 414)

This strategy is similar to the technique of "elevator seeking" on a single disk. The process involves gathering a batch of read requests and sorting them into the most efficient order before performing them. This work is performed by Opus's smart hard disk drive controller on the motherboard, which has its own 16-bit processor and DMA channel

Datasafe's hardware and software security features are effective, and I believe that they can ensure that your PCs are used for their stated purpose and for nothing else.

and employs 304K bytes of extended memory as a disk cache.

You can alter the cache size by using the BIOS utility. By performing DMA transfers to and from the cache memory (which imposes no load on the main CPU), the controller achieves a much higher than normal data transfer rate. This procedure, in turn, permits the disks to be formatted with a 1-to-1 interleave so that a whole track is read per revolution.

Zero latency also further improves the data rate. This feature lets transfers start on any sector boundary without having to wait for the beginning of a track to come around. The data rate also increases because of highly optimized skewing of the sector positions within and between disk platters.

According to the BYTE disk benchmarks, the end result of all this diligent design is that the Opus disk performs three times faster overall than an IBM AT's hard disk. At 0.2 milliseconds, the average seek time is even faster than the Konan Tentime controller fitted to the Minstrel Workstation that I tested (see "HM Systems' Minstrel Workstation," International Section, November 1989 BYTE), although the track-to-track seek time is slightly slower at 6 ms.

The Opus's data transfer rate, as mea-

sured by the CORE test, is 10 times faster than that of an IBM AT. But what actually matters more than benchmarks is that in real applications, the Datasafe's disk really feels fast, much like the Konan Tentime. Because of this feature, you forget that disk accesses used to be noticeable in the bad old days.

One factor that benchmarks cannot bring out is that Datasafe's twin disks improve the speed of every disk read, including the first, whereas systems that rely solely on a cache can only improve the speed of the second and subsequent reads of a data item. In real life, many data items are read just once and so are not sped up at all by a cache system.

The Tseng Lab's EVA/800 VGA card is no faster on text manipulation than an AT, but, as you might expect, it is twice as fast on graphics drawing. Of course, the card's real performance advantage with hardware windowing and scrolling is not measured by the BYTE benchmarks, and it could only be measured by using a program such as AutoCAD, both with and without the appropriate driver. Lack of time precluded this experiment.

Progress or Backsliding?

It's a sad reflection on the state of our society that, after many years when the aim was to make computers easier to use, the Opus Datasafe represents, in a sense, a trend toward making them more difficult to use (or abuse). The hardware and software security features of Datasafe are extremely effective, and I believe that they can ensure that your PCs are used for their stated purpose and for nothing else.

The attraction for military, police, and other security sites is obvious, but, after several public virus scares, many large civilian offices also may now become more concerned with security. The Certus software showed a slightly worrying tendency to fail too safe and lock you out if anything went wrong. Perhaps this is the price of vigilance.

Certus will not appeal to single users. But even without this software, you may still find good reasons to choose the Datasafe. It is a very classy AT clone with excellent disk performance, and when you add certain popular packages, it has the potential for enhanced video performance. Many users with valuable data to protect may welcome the added peace of mind provided by disk shadowing. ∎

Dick Pountain is a BYTE consulting editor, technical author, and software consultant living in London, England. You can contact him on BIX as "dickp."

The HASP Family:
Software protection you can trust.
Software protection you can afford.

As a software producer, you can't market your software without protecting it. Aladdin Knowledge Systems is a leading company in the field of software protection: during the last five years we have enabled hundreds of software producers in more than 30 countries to protect their software.

HASP-3™ connects to the parallel port* of PC/XT/AT and PS/2 computers and compatibles. HASP-3's advanced technology prevents reverse engineering, making the plug virtually un-crackable.

- **Access Password** - A unique password supplied to the software developer is needed in order to access the plug's code. The number of available codes and passwords is enormous (more than 256,000 billion!).

- **The Highest Compatibility** - The plug is transparent to the operation of the computer and the printer.

- **Full Software Support** - All the software required to link HASP-3 to all high-level languages - including the most sophisticated anti-debugging protection and a utility program to protect EXE/COM files - is supplied.

- **Daisy-chaining** plugs can be connected one behind the other.

- **Automatic Virus Detection** can be added to the protected program.

* A serial RS232 HASP is available as well.

MemoHASP™ is the last word in software protection. In addition to all of HASP-3's advantages, **MEMOHASP's 112 bytes or 496 bytes of read/write memory enable you to:**

- Assign a unique code to every software user.
- Control access to different software modules or different software packages.
- Distribute demos which can be activated only a certain number of times.
- Rent or lease software.
- Save passwords or other sensitive information.

Most important of all: all the above operations are performed on the PC without any special programming equipment!
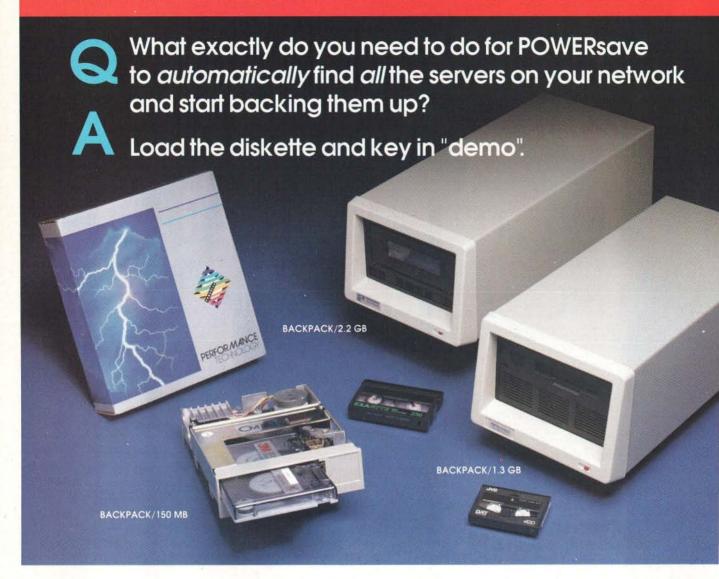
- **There is no Battery in the Plugs!**

For further information contact us at:

ALADDIN
K N O W L E D G E
S Y S T E M S   L T D

130 Dizengoff St., Tel-Aviv
P.O.Box:11141 Tel-Aviv 61110, Israel
Tel: 972-3-226286/241603
Tlx: 35770/1 COIN IL Ext. JBL
Fax: 972-3-243540

© 1985-1990 Aladdin Knowledge Systems Ltd.

Circle 404 on Reader Service Card

SHORT TAKES

BYTE editors' hands-on views of new and developing products

A Pocket PC That Can Exchange Data with a Desktop PC

The Pocket PC from Distributed Information Processing has been discussed in BYTE under its alias, the Atari Portfolio (see Short Takes, August 1989). DIP designed the machine for Atari but retained the right to sell it under its own name in the U.K. In its **Pocket PC Professional** package, DIP has bundled the Pocket PC with a peripheral interface that lets you exchange data with a desktop personal computer.

The Pocket PC Professional comes in a plastic briefcase into which the individual components fit. The components include the Pocket PC itself, three spare AA alkaline cells, a plug-in parallel interface module, a wall-plug transformer for powering the machine from an electrical outlet, a 6-foot cable terminating in 25-pin male D-type plugs, an empty 64K-byte RAM storage card and its battery, a ROM card containing a tutorial and File Manager, DOS software on 5¼- and 3½-inch floppy disks, and several manuals.

For a thorough look at the Pocket PC itself, I refer you to the above-mentioned Short Take. In brief, it is a hand-held device about the size of a child's paintbox (8 by 1 by 4 inches) with an 8-row by 40-column LCD and a miniature QWERTY keyboard. The Pocket PC runs a ROM-resident operating system, which is largely compatible with DOS and uses credit card–size RAM and ROM cards as floppy disks. A set of application programs (i.e., a text editor, a calendar, a diary, an address book, and a spreadsheet) is built into the machine's ROM. After a short period of adjustment, I was pleasantly surprised at how fast I could type on the minuscule keyboard.

It's debatable how DOS-compatible the Pocket PC really is. Few modern DOS programs will run in the 78K bytes of free RAM, and programs that use direct video access run up against the nonstandard video timing. For the programs that do run, the small screen format renders their output unviewable.

The ROM card contains a tutorial, consisting of automated demonstration programs of the built-in applications, and the File Manager. This latter program displays a list of the files in a directory and lets you rename, delete, copy, type, or execute them by pointing with the cursor or tagging groups of files. You can load it and leave it resident so that you return to it instead of DOS after exiting an application. Given the amount of typing it saves and the negligible amount of memory it occupies, it's a pity it wasn't included in the internal ROM.

When I removed the cover on the right side of the Pocket PC case, an edge connector that extends the system bus was revealed. Onto this, you can plug the parallel interface module, a plastic box measuring 3 by 4 inches and bearing a 25-pin D-type socket. The interface is powered from the Pocket PC's own batteries. If you run the cable from this box to the parallel printer port of an IBM PC–compatible desktop, you are ready to communicate. The interface also lets you print directly to a parallel printer using a standard printer cable.

The supplied floppy disks contain among their utilities a file transfer program called FT.COM for the desktop computer; the setup program in the Pocket PC's ROM contains a matching file transfer option. These file transfer programs are very simple in concept, offering just the options Transmit, Receive, Server, List, and Quit. You set either end of the link to Server and then control communication from the other end; usually the Pocket PC will be the Server, and you'll run things from the keyboard of the desktop PC.

I found file transfers easy and reliable, if rather slow; despite the parallel connection, the software uses a proprietary error-correcting serial protocol that limits the data transfer rate to around 2000 bps. The rudimentary user interface was slightly disappointing; something more visual, like LapLink, could surely have been grafted onto the File Manager. The Pocket PC's editor produces plain ASCII text files, which transfer easily into any program on the desktop machine without conversion. The Pocket PC's spreadsheet produces .WKS files that you can read straight into Lotus 1-2-3 and compatible programs.

Another utility program found on the floppy disk is a version of the Pocket PC's built-in applications, called APP.EXE, that runs on the desktop PC so you can keep your diary and address book on both machines. Although it eliminates conversion problems, APP (which is not pop-up) is, frankly, poor competition for PC organizers like SideKick. Other programs on the disk include the standard DOS utilities XCOPY, FIND, SORT, and ANSI.SYS; a print spooler; and a TSR program that simulates keyboard input for use in batch files.

I also received an application called Pocket Finance, which is not included in the Professional kit. This ROM card contains a sophisticated calculator program that can perform scores of financial and statistical computations, like compound interest, amortization, cash flow, break-even analysis, and forecasting with linear regression. In addition, it can draw graphs of the results on the Pocket PC's bit-mapped display.

I found the Pocket PC far more useful than other pocket organizers, because I could type quite lengthy notes on it. While it can't replace a laptop, I could carry it on occasions when a laptop would be an inconvenience. The DOS compatibility angle is something of a red herring, however, because a lot of the popular DOS software won't run satisfactorily on the Pocket PC.

—Dick Pountain
continued

A Full-Featured Shareware Word Processor

You may be surprised to see a shareware product covered in BYTE, but, as PC-Write has demonstrated for several years, shareware programs can be of a professional quality. **EasyEdit II** is a powerful text editor that works with plain ASCII files. It incorporates all the features of a simple word processor, including laser printer and network support, plus advanced editing facilities, such as abbreviations and macros. In addition, you can configure EasyEdit to emulate other word processors.

My copy of EasyEdit II came on a single high-density 5¼-inch floppy disk, but you can get other disk formats or download it from various on-line systems as a self-unpacking ARCed file. To install it, I simply copied the files onto my hard disk, and it ran immediately. EasyEdit II reads in from nine auxiliary files that hold default settings, abbreviations, macros, mouse support, menus, matching pairs, and key assignments; on top of help, manual, and printer font files, this makes quite a full directory. However, you don't need to worry about any of these files at first; furthermore, EasyEdit II generates most of them for you.

For example, to change the screen colors, you just pull down the Options menu and choose colors from a Side-Kick-like color chart. When you save the options, EasyEdit automatically generates a color file called EECOLORF .ILE.

Most of the auxiliary files are plain ASCII text, so you can edit them with EasyEdit itself. You also write colors and key bindings into the EE.EXE file, so it will always work sensibly without auxiliary files. The auxiliary files come into their own on a network where you can keep EasyEdit in a directory on the server with some global defaults, but each user can also have a local di-

rectory with his or her own customized colors, macros, and so on; two DOS environment variables called EE-PATH and EELPATH point to these directories. EasyEdit II does not use the file-locking capabilities of DOS 3.0 but has its own mechanism, which works under DESQview, too.

The program first presents a clean editing screen with just a status line at the top showing the current line, column, and page and the state of various toggles. To select commands, you press the F10 key, which brings up a menu bar with pull-down menus. The user interface resembles Borland's Turbo Pascal editor, with interaction via pop-up dialog boxes or point-and-shoot directory windows rather than a command line. You can also control the editor with key sequences, which, by default, are WordStar-compatible, but you can change them using a separate interactive key editor called KBD_INST. The extensive indexed on-line help system is automatically updated to reflect your new assignments—a neat feature. There are 190 possible editor actions to assign, and I found a subset that emulates all my preferred key sequences from PC-Write. It's even possible to customize the menu system itself by simply editing a text file of plain English descriptions.

A memory-resident editor,

EasyEdit II cannot edit files bigger than the amount of available RAM, but it can use all of DOS's 640K bytes; I had 300K bytes to play with. You can load many files simultaneously in Sprint-like horizontal windows. Also like Sprint, EasyEdit lets you save to disk at timed intervals, and it maintains a .SAV file that you can use to recover your text in the event of power cuts or other disasters. When you boot up EasyEdit, it remembers the last file you edited and returns you to the exact row and column where you left off.

EasyEdit II features word wrapping and gives you the option of automatic space compression or justification, but it lacks an automatic paragraph reformat. I prefer the old Control-B-style reformat, but some people will sorely miss this feature. The program caters to professional writers with word-count and table-of-contents-generation facilities, but there are no facilities for footnotes or outlining. There is an elegant tab mechanism for columnar material—you can set tabs with a single keystroke at the start of all words on the current line, producing an instant table. You can also embed tab lines in a file to vary the margin settings. Printer control is very simple, supporting just the usual attributes of boldface, underline, italic, superscript, and so on. However, you can change font

files on a laser printer by including the equivalent of dot commands.

The search-and-replace facilities are EasyEdit's weakest feature—WordStar-like, but with no wild cards or special line-end symbols (you have to use Control-P, Control-M, Control-P, and Control-J). There is, however, a handy option that lets you apply a named macro to each found text occurrence. The Super-Key-look-alike macro editor, which assigns macros to the Alt/1–9 keys, is one of Easy-Edit's best features; you can also pick macros by name from a menu window. Macros can include "live" dates, times, and filenames and can prompt you for input.

The Abbreviation mechanism lets you set up a file of two- or three-letter abbreviations, which you can expand with just one keystroke. There is no built-in spelling checker, but the hooks are there (including a menu option) for using Borland's Turbo Lightning engine. The program's Matched Pairs mechanism is designed for programmers; it allows you to define a file of pairs like BEGIN...END or {...} so that EasyEdit can become a syntax-smart program editor.

EasyEdit also offers simple but elegant ways to manipulate columns and draw boxes. It has no graphics drawing facilities, but you can append binary files to a text file (and extract them) for E-mail.

EasyEdit II is easy enough for a beginner when used through menus, but its real appeal is to experienced word processor users who feel jaded with the "quasi-desktop publishing" complexity of modern big-name word processors and are looking for speed, simplicity, and a plain ASCII format. In addition, its almost unparalleled degree of configurability lets you, in effect, design your own editor.

—*Dick Pountain*

THE FACTS

EasyEdit II 3.21
Shareware with one month free use;
£20 registration fee for a printed manual and one free upgrade; £20 each for two additional upgrades; company site licenses start from £160 for 10 users.

Requirements:
An IBM PC with 256K bytes of RAM, DOS 2.0, and one floppy disk drive.

The AsEditCo
Warren Cottage
Warren Lane
Elmswell
Suffolk IP30 9DT
UK
44-0359-41601
Inquiry 884.

USE YOUR FAX TO FIND THE LOWEST COMPUTER PRICES IN THE WORLD

When you know exactly what you want, the last thing in the world you need is a salesman.

Once you've come to a thoughtful decision on components, systems and peripherals or software, all you want to know is "how soon can I get it."

...and you want the facts "now!"

That is why we formed Elex International five years ago. We don't want to sell you a thing. We do want to help you find exactly what you need at a price you are happy to pay.

Very often, we find what you want in the very last place you might expect: in the USA.

Unrestricted competition in the United States keeps prices low and availability high.

Of course, many of the same powerful systems, soft-ware programs and components are available to some extent throughout the world, but multi-tiered international distribution channels can inflate prices well above comparable levels we find in the US.

So, whether you are looking for anything from a single 41256 memory chip, to a fully configured IBM AS400 minicomputer system, Elex International gives you an important new purchasing option.

We provide this service to end users, value added resellers and original equipment manufactures throughout the free world.

Wherever you are—in Europe, Africa or Asia—you can get direct access to the low prices and wide range of choices available in the US, right now.

It's easy to test the Elex solution. Fax your system requirements, detailing all appropiate specifications, to an Elex order center. We will respond with a firm bid within the span of a single business day.

Fax any requirements written in English directly to our US Order Center for immediate response.

Requirements written in French, German, Spanish Italian or Russian should be Faxed to our Order Center in Geneva.

If you don't need an immediate quote, feel free to call or write at any of our Order Centers.

However you choose to contact us, you will gain an important new source that is committed to filling your computer system requirements promptly and at prices your local distributor cannot match.

Payment Method: L/C, T/T, VISA, COD

Elex International, Inc.

136 North 3rd St.
Philadelphia, PA 19106, USA
Telephone: 215-627-1691
Toll-free Orders: 800-783-2313
FAX: 215-627-2342

Elex Information Systems, SA

16 Rue du Simplon
1207 Geneva, Switzerland
Telephone: 41 22 7864971
24-hour Elex Net Modem 41 22 7357280
FAX: 41 22 7357669

Drawing for Microsoft Excel Spreadsheets

DrawCel is intelligent drawing software for Microsoft Excel that lets you link the shape and location of drawings to user-defined spreadsheet references. If relationships or dimensions change, the drawings will automatically reflect these changes.

The drawings are composed of basic geometric elements (i.e., lines, rectangles, circles, ellipses, polygons, and text), just as in traditional drawing programs. The drawing functions define position and size, filling pattern, line width, and rotational angle of the respective object. In this way, you can enhance drawings and graphics with text at any angle.

You define the drawing by means of one or more spreadsheets, with every object based on a drawing function of the sheet. You place the drawing functions in a cell and define the arguments as references to other cells. These arguments govern, depending on the type of function, qualities such as the height, width, and radius of the objects.

After you have defined the drawing functions in a spreadsheet, you build up the actual drawing in a separate window using a key combination or by selecting a command in the extended Excel menu bar.

The graphics window is equipped with a zoom function and commands for printing and copying the drawing. The printing options are Excel standard; DrawCel supports all Microsoft Windows printers. With the copy command, you can transfer the drawing to the clipboard and use it in other Windows applications. By clicking on a sensitive area in the DrawCel graphic with a mouse, you can

With DrawCel drawing software and Microsoft Excel spreadsheets, you can link the shape and location of drawings to user-defined spreadsheet references.

load predefined Excel spreadsheets or execute Excel macros.

DrawCel requires a 286- or 386-based computer with EMS and the latest versions of Microsoft Excel and Microsoft Windows.
Price: 1662 deutsche marks.
Contact: Management Information Systeme GmbH, Bunsenstrasse 22, 6100 Darmstadt, West Germany, 49-6151-8-20-77.
Inquiry 944.

Two 386SXes from Elonex

Elonex offers two computers with 386SX processors; one is a laptop, and the other is a desktop model.

The LT386SX laptop weighs just 6.5 kg and comes with a 16-MHz CPU, 1 MB of on-board single-in-line-

memory-module RAM (expandable to 4 MB), a 640- by 480-pixel VGA display, a 3½-inch 1.44-MB floppy disk drive, and a 40-MB hard disk drive. The laptop also comes with Microsoft Windows/386, Microsoft Works, DOS 4.01, and a carrying case with shoulder strap.

The keyboard has 85 keys with tactile feedback; an external IBM AT–compatible keyboard connector is available as an option. LED indicators show you the status of power, the floppy disk drive, the hard disk drive, Caps Lock, Num Lock, and Scroll Lock.

The VGA display is a gas-plasma panel with an active display area of 211 by 158.4 mm. The screen offers a 120-degree viewing angle, a contrast ratio of 1 to 10, and 16 gray scales.

Options for the LT386SX include an 80387SX coprocessor and an external expan-

sion box providing two AT slots and one XT slot.

The Elonex 386SX/B is a compact desktop computer based on Intel's 16-MHz 386SX microprocessor. It features a motherboard with one Centronics parallel port, two RS-232C serial communications ports, a built-in VGA display adapter, a floppy disk drive controller capable of handling two drives, and an IBM AT bus for a hard disk drive with an embedded controller. In addition, the machine offers five expansion slots and up to four half-height storage devices.

An entry-level Elonex 386SX/B includes 1 MB of memory, a 3½-inch 1.44-MB floppy disk drive, a 3½-inch 40-MB hard disk drive, a tilt-and-swivel VGA monochrome monitor, a 102-key keyboard, and a Microsoft-compatible mouse. Also included with the machine are DOS 4.01, Microsoft Windows/386, and Microsoft Works.

You can expand the system memory on the CPU board up to 8 MB with built-in support for EMS 4.0; you can expand the system memory to a total of 16 MB using a memory-expansion card. Options for the Elonex 386SX/B include a tilt-and-swivel color VGA monitor, 100- and 200-MB hard disk drives, and Ethernet and multiport cards for Novell NetWare-based LANs and Unix-based multiuser systems. You can use a diskless version of the entry-level machine as a LAN workstation. It comes equipped with a NetWare boot ROM and a 10-Mbps Ethernet connector.
Price: £1795 for the LT386SX; £1345 and up for the 386SX/B.
Contact: Elonex plc., Rays House, North Circular Rd., London NW10 7XB, UK, 44-1-965-3225.
Inquiry 916.

If you would like your new product considered for publication in the international section of BYTE, send press releases to BYTE, Attention: Martha Hicks, One Phoenix Mill Lane, Peterborough, NH 03458, U.S.A.; or Dick Pountain, BYTE, McGraw-Hill Publishing Co., 34 Dover St., London W1X 3RA, UK; or Nikkei BYTE, 1-1, Kanda-Ogawamachi, Chiyoda-ku, Tokyo 101, Japan. All press releases must contain price information, address, and telephone number.

Choose Your Serial Channel

You can set each of the eight serial channels on the MSC-8 multichannel serial card to RS-232C, RS-422, RS-485, or 20-mA current loop at any individual address location and any interrupt channel, according to Blue Chip Technology. The card is designed for applications such as communicating with large numbers of intelligent sensors and instruments, LANs, shop-floor data collection, and time and attendance systems.

You select the RS-232C, RS-422, RS-485, or 20-mA current loop operation using small submodules that plug into the board. Other channel-by-channel options include the base address (128 possible addresses per channel) and interrupt channel (IRQ2/7), which you set via DIP switches.
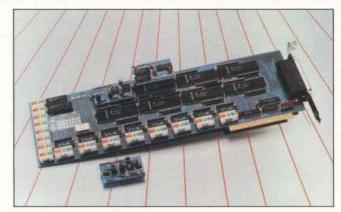
The card's software driver lets you configure the board rate via the IO Control Channel and gives you flexible addressing via DOS from COM1 to COM99.

You can use the MSC-8 in an IBM PC, XT, AT, or 386 computer that has a standard 8-bit PC-bus interface. You can install multiple cards in a single PC up to the maximum number of available slots.
Price: £195 (two-channel) to £495 (eight-channel).
Contact: Blue Chip Technology, Hawarden Industrial Park, Manor Lane, Deeside, Clwyd CH5 3PP, UK, 44-244-520222.
Inquiry 943.

Using small submodules that plug into the MSC-8 multichannel serial card, you can select RS-232C, RS-422, RS-485, or 20-mA current loop operation.

A SQL Interface for Common Lisp Programming

With LispWorks' new Structured Query Language (SQL) interface, you can write Common Lisp applications that can simultaneously access different databases. For example, a program could retrieve data from an Oracle database, process it, and then store it in an Informix database.

The interface employs object-oriented tools and techniques that let you retrieve, insert, modify, and delete data from a number of relational databases in a familiar Lisp environment. It supports SQL cursors and provides a generic routine that takes SQL update statements.

LispWorks includes the standard object-oriented extension of Lisp, known as Common Lisp Object System (CLOS), and the standard toolkit package, Common Lisp User Environment (CLUE). The LispWorks Interface Toolkit, built on CLUE, contains components such as buttons, menus, and scroll bars that help you design your own user interfaces.

LispWorks runs under the X Window System and supports multiprocessing on Sun-3, Sun-4, DECstation, Sun386i, and SPARCStation 1 machines. In addition to SQL, you can also interface LispWorks to code written in C or FORTRAN.
Price: $4000 to $5000 U.S. for a single license, depending on the machine.
Contact: Harlequin Ltd., Barrington Hall, Barrington, Cambridge CB2 5RG, UK, 44-223-872522.
Inquiry 901.

A High-Speed Hard Disk Drive with Removable Disks

The ProDrive 40 for the Macintosh combines a 40-MB hard disk drive with the flexibility of a removable floppy disk. You get a data transfer rate of 10 Mbps and an average access time of 25 ms. User-friendly software lets you use password protection to guard important data and provides simple and fast operation. The ProDrive 40 measures 25 by 25 by 6.5 cm and weighs 4.5 kg.

Formac also offers the ProTape 60 for the Mac, a tape backup unit that gives you a backup rate of more than 3.3 MB per minute and lets you back up your entire drive (file by file) with the click of a button. The unit automatically backs up hard disks of more than 60 MB over two or more tapes; the tape capacity is 50 or 60 MB. Advanced error-correction algorithms ensure that your data is backed up and restored accurately. You attach ProTape to your Mac's SCSI connection. It measures 25 by 25 by 6.5 cm and weighs 3.5 kg.
Price: 3300 Deutsche marks each.
Contact: Formac GmbH, Charlottenstrasse 13, D-1000 Berlin 61, West Germany, 49-30-251-04-01.
Inquiry 929.

Landscape Printing on a LaserJet II

Now you can print text documents in landscape mode on a Hewlett-Packard LaserJet II printer with Codespeed's A4List utility. The utility lets you print text documents on A4 paper—two pages on each side in double-sided mode, or two pages on one side in single-sided mode. (For double-sided printing, you have to pass the paper through the printer twice.)

Because the utility uses an internal line printer font, it is necessary to download a font to the printer. You can set the number of lines per page at 62, 64, or 66, and the number of characters per line at 80 or 93. Suitable margins are maintained for installing the output in a standard A4 folder. A4List automatically places a title, date, time of printing, and page number at the top of every page.

A4List runs on the IBM PC, XT, AT, or PS/2s.
Price: £25.
Contact: Codespeed, Codespeed House, 21 Sewell Close, St. Albans, Hertfordshire AL4 0TB, UK, 44-727-37488.
Inquiry 919.

Give Your Computer a Voice

HIFIGEN is a range of peripherals for recording, editing, and replaying high-quality sound using your computer's hard disk drive as sound memory. The system brings high-quality sound dimension to programs, communication scenarios, and computer-aided teaching programs, with commands driven by software such as Concorde, Show Partner, and Storyboard.

The system uses a split-band adaptive differential pulse-coded modulation (ADPCM) coding technology that offers high-quality sound for a relatively low data rate. During the coding, the system digitizes sound, and compression algorithms then reduce the amount of information necessary to represent it. HIFIGEN stores the end result on your hard disk in compressed form.

To replay the sound, the system applies a reverse process—extracting information from the hard disk, expanding it, and then converting it into an analog signal that you can replay.

With Sound Studio, you can add sound to your programs and scenarios. It lets you create and edit sound files on your hard disk from a tape recorder, a compact disk player, or a microphone.

HIFIGEN-REC is a board for the IBM AT that lets you record, replay, and edit sound files. You input the sound files from a tape recorder, microphone, or disk player and output them to headphones, a preamplifier, or a 2-W loudspeaker.

With the HIFIGEN-EDIT sound-editor software, you can record sound files, cut and paste parts of sound sequences, and put marks for later synchronization of pro-

The HIFIGEN range of peripherals lets you record, edit, and replay high-quality sound using your computer's hard disk drive as sound memory.

grams with sound sequences.

HIFIGEN-PC is an IBM PC board that lets you replay sound files to headphones, a preamplifier, or a 2-W loudspeaker.

You connect HIFIGEN-BOX to your computer's parallel printer port. It is a smart stand-alone and laptop baffle that is a real "sound printer" with the same function as the HIFIGEN-PC board, except it is simpler to connect and disconnect via the parallel port. The baffle, with its integrated amplifier in the box, has been designed especially to use the characteristics of the digital sound—output goes to headphones or a preamplifier.

Every HIFIGEN product comes with a software package that includes basic sound replay commands, such as replay a sound file, stop playing a sound file, wait for the end of a sound file, wait for the next synchronization mark in a sound file, set output audio level, and start or stop playing (de)crescendo.

You can access the commands at the DOS command level or through software interrupts. The products perform sound-playing functions through interrupts, which lets you run other tasks simultaneously.

The HIFIGEN products require an IBM AT with DOS 3.0 and a 30-MB hard disk drive, which can hold more than 1 hour of high-quality 7-kHz bandwidth sound.
Price: 9700 French francs each for HIFIGEN-REC and HIFIGEN-PC; 15,100 FF for HIFIGEN-BOX.
Contact: Systèmes Grenat, Z.I. Nord, Route de Trégastel, 22300 Lannion, France, 33-96-48-85-18.
Inquiry 926.

Accounts II for Unix

Accounts II Top End for Unix provides its own window management, configurable color, and context-sensitive help screens. The package consists of modules for Sales Ledger, Purchase Ledger, Nominal Ledger, Payroll, Stock Control, Invoicing, Purchase Order Control, Sales Order Processing, Names and Addresses, Job Costing, Cashbook, Fixed Assets Register, Spreadsheet, and Desktop utilities. It also includes multicurrency capability, a report generator, and bank automated credit service functions.

You can access data files you have created with Accounts II with database packages like Ingres and Oracle that provide Structured Query

Language support. Thus, users who are already familiar with Unix database packages can access Accounts II data files for inquiry and reporting.

The Spreadsheet module is compatible with Lotus 1-2-3 and provides direct links to the accounting data files. You can use it for management reporting and business graphics.

The first version of Accounts II Top End is available for IBM's 6150 and PS/2 Models 70 and 80, all running under AIX.
Price: £375 to £800 per module.
Contact: Graffcom Systems Ltd., 8 Church St., Old Isleworth, Middlesex TW7 6BG, UK, 44-1-847-4604.
Inquiry 909.

Transfer Files Between Unix and IBM Mainframes

You can transfer files from a Unix computer to an IBM mainframe with INCAA Datacom's combination hardware and software file transfer package. The system handles all the IBM protocol emulation outside the Unix machine with a microcomputer protocol converter, called the Programmable Interface Translator, which you connect to your Unix computer through an RS-232C serial port. You can transfer data at any speed up to 19,200 bps, and the software lets you transmit and receive files in BSC 2780-3780 or SNA/SDLC 3770 format. The operating system automatically handles error and status messages.
Price: About £2000.
Contact: INCAA Datacom Ltd., 75 High St., Harrold, Bedford MK43 7BJ, UK, 44-234-720144.
Inquiry 903.

5 out of 5 hackers prefer other software protection methods to Hardlock E-Y-E.®

What hackers dislike...

Hardlock E-Y-E was designed using cryptographic principles. It took the experience and know-how of Germany's No. 1 in software protection and the leading edge technology of a US semiconductor company to create the ultimate software protection tool. Hardlock E-Y-E is based on a custom chip featuring secure algorithmic response rather than simple bit swapping or counting schemes.

What software developers like...

Hardlock E-Y-E combines all the features software developers require in a single product: algorithmic response to provide security and an optional non-volatile memory to allow custom configurations. FAST Electronic has made implementation of Hardlock E-Y-E in your software easy. Use HL-Crypt to protect .EXE or .COM files, or incorporate high level language interface routines in your software. The algorithm parameters and the contents of the memory can be programmed in seconds using our Crypto-Programmer card. This unique card guarantees that no one else can burn your original codes. Simply plug the card into any PC slot and start up your own Hardlock E-Y-E workshop.

What your customers will like...

Hardlock E-Y-E allows unlimited backup copies. The device is shipped with the software for the user simply to plug into the parallel interface and forget.

Daisy chainability, outstanding reliability (no battery is needed), and the most compact High-Tech design ensure that your customer will accept Hardlock E-Y-E.

What your accountant will like...

Hardlock E-Y-E needs no factory coding. This ensures optimum delivery schedules and stock flexibility. Revenues will go up as software piratry and multiple usage are prevented. Despite its wealth of features, Hardlock E-Y-E's prices remain competitive.

...As more and more software developers, customers and accountants appreciate the Hardlock E-Y-E device, hackers like it less and less.

Hardlock E-Y-E
programmable, algorithmic response and memory option – all in one.

Organize Your Video Images

VIM 1.01, the Video Manager, can overlay text on video recordings for documentation purposes, correlate and compare two video images by calculating the difference or contrast between them, and reduce electronic noise in an image by averaging several frames. In addition, VIM performs two-level correlation between live video and a stored image in real time, stores and reads images to and from disks, and gives you access to an image-processing interpreter.

The image-processing software efficiently organizes and correlates TV camera and videocassette or tape recorder image data, using a personal computer and a video frame-grabber board called PIP-1024. VIM uniquely marks and documents video images and then retrieves them from videocassettes, tapes, and disks. The system also lets you perform image surveillance, improvement, and control, in either real time or time lapsed from videotape recordings.

The menu-driven software consists of six pull-down submenus. The Mark submenu lets you mark video images before you record them on a VTR or VCR. The Grab submenu grabs images directly from TV cameras, VTRs, VCRs, and other video sources. If you have marked the images with VIM, the software can grab the images using a unique frame number or event number, or you can grab images by simply pressing a key.

The Correlate submenu correlates and compares two video images by calculating the difference or contrast between them. You can display the resulting image as a bipolar or a unipolar image. To get optimum performance, you can adjust the bipolar and unipolar contrast of images. You can also write comments on the resulting image or on other images that you select.

The Average submenu reduces electronic and optical noise by averaging any number of images. You add each new image to the earlier ones with one keystroke.

With the On-line submenu, you can observe real-time correlation images by combining the live video with a stored reference image. Both the reference image and the live video are divided into two gray levels before they are logically combined. You can update the reference image at any time and display the resulting images as live video on the video monitor. VIM lets you store up to three extra images with full dynamic range (256 gray levels) for later retrieval, correlation, and analysis.

The Utilities submenu lets you access the PIP interpreter program without leaving PIP. You can also view, save to disk, and read from disk images you have stored in PIP-1024, as well as select TV standards and adjust the analog input signals from software.

VIM 1.01 requires a TV camera, VCR, or VTR producing video images according to RS-170/RS-330 CCIR TV standards; a video monitor accepting video images according to the same standard; and an IBM PC with 384K bytes of RAM, DOS 2.1, and a Matrox or PIP-1024 frame-grabber board; a hard disk drive is recommended.
Price: 10,500 Norwegian kroner.
Contact: Conoptica AS, Finnmyrveien 2A, N-7060 Klaebu, Norway, 47-7-83-24-51.
Inquiry 925.

continued

Compuclassics

We Rubbed Our Magic Lamp, And Poof!! Magic Prices

PLEASE CALL US FOR PRODUCTS NOT LISTED. WE SHIP TO APO & FPO ADDRESSES. CALL US FOR A FULL CATALOG.

3.5" Disk Versions May Cost More

ACCOUNTING & PERSONAL FINANCE

Andrew Tobias Tax Cut 1989	49.00
Dac Accounting 4.0	89.00
Dac Bonus Pack 4.0	175.00
Managing Your Money 6.0	119.00
Money Counts	35.00
Peachtree III w/ Data Query	229.00
Personal Lawyer	39.00
Quicken 3.0	39.00
Ronstadt's Financials	179.00
Timeslips III 3.4	169.00
Turbo Tax Personal	45.00
Turbo Tax Personal State	29.00
Turbo Tax Professional	229.00
Turbo Tax Professional State	169.00
Willmaker 3.0	35.00

CAD

Autosketch 2.0	95.00
Design Cad 3D 2.2	205.00
Design Cad 4.0	155.00
Drafix Cad Ultra	259.00
Easy Cad 2	132.00
Generic 3D Drafting	140.00
Generic Cadd Level 3	165.00

COMMUNICATIONS

Carbon Copy Plus 5.2	115.00
Close Up Lan (8 User)	529.00
Close Up Support 3.0	165.00
Close-Up Customer 3.0	129.00
Crosstalk Mark 4	139.00
PC Anywhere III	69.00
Procomm Plus	50.00
Reflection 2 PLus	199.00
Reflection 4 PLus	295.00
Relay Gold	175.00
Smarterm 240 3.0	199.00
Smarterm 320	115.00

DATABASE

Clarion Professional Developer	399.00
Clipper (Summer '87)	429.00
Dataease 4.2	485.00
DB Fast/DOS Plus	179.00
DBASE IV	499.00
Foxbase LAN	299.00
Foxbase Plus	199.00
Foxbase S/U Runtime	265.00
Foxpro Lan	679.00
Foxpro Single User	479.00
Genifer	209.00
Paradox 3.0	469.00
PC File 5.0	79.00
PFS Professional File 2.01	199.00
R & R Code Generator	109.00
R & R Relational Report Writer	109.00
RBase For DOS	465.00
UI Programmer	349.00

DESKTOP PUBLISHING

Bitstream Fonts (each)	95.00
Dan Bricklin's Page Garden	65.00
Finesse	95.00
Freedom Of Press	269.00
Go Script 3.0	99.00
Go Script Plus 3.0	195.00
PFS 1st Publisher	89.00
Power Up! Express Publisher	89.00
Printmaster Plus 2.0	35.00
Publish It!	119.00
Springboard Publisher	79.00
The New Print Shop	39.00
Ultrascript PC 2.0	119.00
Ultrascript PC Plus 2.0	269.00
Ventura Professional Extension	379.00
Ventura Publisher	549.00

EDUCATION & ENTERTAINMENT

ATI Teach Yourself Series (EA)	59.00
F-15 Strike Eagle II	35.00
F19 Stealth Fighter	45.00
Flight Simulator 4.0	45.00
Individual Training Series (EA)	40.00
King's Quest (EA)	35.00
Leisure Suit Larry III	39.00
MS Learning DOS	35.00
Simcity	35.00
Space Quest III	39.00
Test Drive II	27.00
Tetris	22.00
Typing Tutor IV Plus	35.00
Vette!	32.00
Where ... Is Carmen San Diego	30.00

FORM DESIGN & FLOWCHARTING

Corporate Ladder	49.00
Flowcharting II Plus	135.00
Formtool W/Greatest Hits 2.01	55.00
Formworx with Fill File 2.5	85.00
Interactive Easyflow 6.1	109.00
Jetform Design	259.00
Org Plus Advanced	69.00
Per:FORM 2.0	169.00

GRAPHICS

Applause II	65.00
Autodesk Animator	249.00
Colorix VGA Paint	105.00
Deluxe Paint II Enhanced	89.00
Drawperfect	285.00
Freelance Plus 3.01	349.00
GEM/3 Draw Plus	179.00
Graph-In-The-Box Executive	185.00
Grapher	145.00
Graphwriter II	319.00
Grasp by Paul Mace	119.00
Harvard Graphics 2.13	319.00
Hijaak	95.00
Hot Shot Graphics	149.00
Inset Plus (w/HiJaak)	139.00
Micrografx Xport	285.00
PC Paintbrush IV	65.00
PC Paintbrush IV Plus	119.00
Pizazz Plus 1.3	69.00
Print A Plot 2.0	199.00
Publisher's Paintbrush	169.00
Show Partner FX 3.5	229.00
Surfer	359.00
Xerox Graph	195.00

HARDWARE & PERIPHERALS

ATI 2400 Etc Modem Int. w/MNP	169.00
ATI VGA Wonder 512 w/Mouse	349.00
Complete Commmunicator	585.00
Complete Fax 9600	459.00
Complete Half-Page Scanner	189.00
Copy II Option Board Deluxe	115.00
Hicard AT w/0K	249.00
Intel Above Board Plus 512K	435.00
Intel Connection Coprocessor	575.00
Kensington Expert Mouse for PS/2	115.00
Logitech Bus Mouse C9	85.00
Logitech Scan ManPlus	185.00
Logitech Serial Mouse C9	60.00
Logitech Trackman Serial or Bus	89.00
Masterpiece	89.00
Masterpiece Plus	105.00
MS Mouse w/Paint (Ser. or Bus)	109.00
MS Mouse w/Windows (Ser. or Bus)	139.00
Microspeed Fasttrap Serial	99.00
Orchid Prodesigner Plus 512K	349.00
Pacific Data 2+2 (2mb)	395.00
Pacific Data 25-in-one	289.00
Pacific Data Pacific Page	509.00

Pacific Data Pacific Page IIP	389.00
Paradise VGA 1024-512	305.00
Polaroid Palette Plus (EGA)	2399.00
Prac Periph 2400 Ext w/MNP	209.00
Prac Periph 2400 Int w/MNP	179.00
Sota 286i Accelerator	259.00
Sota 386i Accelerator	399.00
Sota VGA 16 512K w/Mouse	339.00
Super Cartridge 1 -IQ Engineering	275.00
Super Cartridge 2 -IQ Engineering	449.00
Sysgen Bridge File with Adapter	275.00
Worldport 2400 Modem	249.00
Worldport 2400 Modem w/ MNP	345.00
Worldport 2496 Fax Modem	485.00

INFORMATION MANAGEMENT

Act!	229.00
Agenda	275.00
askSam	175.00
Gofer! 2.0	49.00
Grandview	199.00
Infoselect	79.00
Instant Recall	79.00
IZE	285.00
Memory Mate 3.01	45.00
Who What When 1.09	119.00

INTEGRATED

Alphaworks	115.00
Informix Smartware II	469.00
MS Works 2.0	95.00
PFS First Choice	115.00
Q & A 3.0 (Networkable)	245.00
Symphony Plus	475.00

LANGUAGE & PROGRAMMING

.RT Link	215.00
Brief	159.00
BTrieve	159.00
C Asynch	115.00
C Tools Plus	95.00
Clear + for C (or dBase)	135.00
DGE	125.00
K Edit	115.00
Lattice C Compiler	165.00
Macro Assembler	99.00
Matrix Layout 2.0	135.00
MS Basic 7.0 Development System	329.00
MS C 5.1	299.00
MS Fortran	299.00
MS Quick C	67.00
MS Quick C w/ Quick Assembler 2.01	135.00
MSQuickbasic	67.00
Norton Editor 1.3	45.00
Object Professional	109.00
Qbase	149.00
Quickpac Professional	155.00
Quickpack	72.00
Smalltalk V	75.00
Smalltalk V286	129.00
Sourcer w/Bios	109.00
SPF/PC	159.00
Turbo Asynch +	115.00
Turbo B Tree Filer S/U	79.00
Turbo C	105.00
Turbo C Pro Pack	175.00
Turbo C Tools	95.00
Turbo Pascal 5.5	105.00
Turbo Pascal Pro Pack	175.00
Turbo Power Tools Plus	95.00
Turbo Professional	79.00
Windows Development Toolkit	329.00
Zortech C ++ 2.0	149.00

NETWORKING

Lanspool	289.00
Lantastic 2MBPS Adapter	189.00
Lantastic 2MBPS Starter Kit	409.00
Lantastic Ethernet Adapter	259.00
Lantastic Ethernet Starter Kit 10 MBPS	555.00
PS Print Advanced	255.00
Tops For Dos	115.00
Western Digital Ethercard Plus	179.00
Wordperfect Office	279.00

OPERATING ENVIRONMENT

Desqview 2.26	79.00
Desqview 386 2.2	125.00
IBM DOS 3.3	95.00
IBM DOS 4.01	119.00
IBM OS/2 Standard Edition 1.2	269.00
MS Windows 286	67.00
MS Windows 386	129.00
PC/MOS 386	175.00
Software Carousel 3.01	52.00
VM/386	155.00

OS/2

Brief For O/S2	205.00
Displaywrite 5/2 Composer	359.00
Excel for OS/2	329.00
Logitech Multiscope	185.00
MS OS/2 Presentation Mngr Toolkit	329.00
Multiboot	49.00
Paradox OS/2	509.00
Smalltalk/V PM	329.00
Wordperfect for OS/2	319.00

PROJECT MANAGEMENT

Adv. Project Manager Workbench 3.05	845.00
Harvard Project 3.0	465.00

Scitor Project Scheduler IV	429.00
Superproject Expert	449.00
Timeline 4.0	479.00

SCIENTIFIC / STAT

Brainmaker	149.00
Chiwriter	139.00
Derive	130.00
Mathcad 2.5	305.00
Mathematica386	595.00
Mathematica 386 w/387 Support	845.00
SPSS/PC Plus	549.00
Statgraphics	579.00
Systat w/ Graphics	509.00

SPREADSHEET

Allways (Lotus 1-2-3 or Symphony)	85.00
Lotus 1-2-3 Ver. 2.2	349.00
Lotus 1-2-3 Ver. 3.0	419.00
Lucid 3D	62.00
Planperfect	279.00
Quattro Pro	269.00
Quattro Pro Upgrade	99.00
See More Lotus 2.0	59.00
Sideways 3.2	42.00
Supercalc V	319.00
The Baler	359.00
VP Planner Plus 3D	152.00

UTILITIES

1 Dir Plus 3.0	49.00
386 To The Max	45.00
386 To The Max Pro	85.00
Above Disk 3.1	65.00
Above Lan	285.00
Automenu	39.00
Brooklyn Bridge 3.0	75.00
Certus 1.0	99.00
Check It	89.00
Copy II PC 5.0	25.00
Copywrite/Zerodiskw/Rescue	55.00
Direct Access 5.0	52.00
Disk Technician Advanced	95.00
Disk Technician Pro	42.00
Fastback Plus	109.00
Fastrax	35.00
Headroom	79.00
Kolod Htest/Hformat 3.0	79.00
Lap Link Release III	92.00
Mace Utilities 1990	89.00
Magellan	139.00
Norton Backup	89.00
Norton Commander 3.0	89.00
Norton Utilities Advanced 4.5	79.00
PC Tools Deluxe 6.0	77.00
Pop Drop Plus	65.00
Print Cache (Lasertorq)	105.00
Print Q 4.0	80.00
Q-Dos II	39.00
QEMM 5.0	65.00
QRAM	49.00
Softsafe 2.0	80.00
Software Bridge 4.1	85.00
Speedstor	35.00
Spinrite II	69.00
Super PC Kwik	49.00
Super PC-Kwik Powerpak	79.00
Turbo EMS 5.0	69.00
V Cache	45.00
Word Perfect Library 2.0	75.00
X-Tree Net	232.00
X-Tree Pro	55.00
XTree Pro Gold	75.00

WINDOWS

Adobe Illustrator Windows	295.00
Adobe Streamline	249.00
AMI Professional	319.00
Arts & Letters Graphics Editor 2.0	465.00
Corel Draw Windows 1.1	329.00
Crosstalk For Windows	125.00
DB Fast/Windows	175.00
Excel	319.00
Formworx System 2	175.00
hDC Windows Express	45.00
hDC Windows Manager	45.00
Hyperpad 1.0	95.00
IBM Current	265.00
Micrografx Designer	469.00
Micrografx Windows Graph Plus	355.00
Omnipage 386 2.1	599.00
Pagemaker 3.0	529.00
PC Paintbrush Plus for Windows	99.00
Pubtech File Organizer	159.00
Superbase 4 for Windows	439.00
Word for Windows	329.00
Xerox Formbase	325.00

WORD PROCESSING
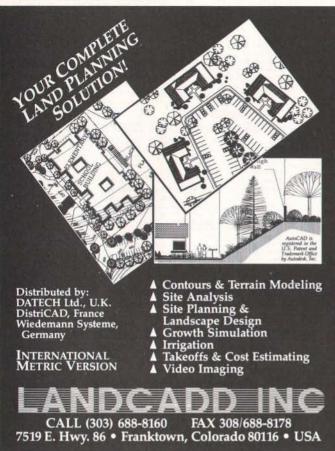
Displaywrite IV	255.00
Grammatik IV	48.00
Multimate 4.0	295.00
Nota Bene	335.00
Rightwriter	52.00
Word 5.0	219.00
Word Perfect 5.1	249.00
Word Perfect Net Add-On 5.1	169.00
Wordstar 5.5	199.00

XENIX/UNIX

SCO Foxbase + 386	675.00
SCO Unix Dev. System 386	675.00
SCO Unix Operating System 386	619.00
SCO Xenix Complete System 386	1079.00
SCO Xenix Operating System 386	479.00
Wordperfect for SCO Xenix 386 (5 user)	519.00
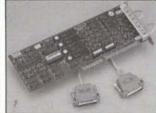
ProNet LAN Products

Proteon offers a new generation of network adapters and an intelligent wire center that supports unshielded twisted-pair cabling over 4- and 16-Mbps IEEE 802.5 and token-ring networks. The products feature unified media support and advanced network management capabilities that simplify networking installations.

The new p1308 Network Interface card for Proteon's ProNet-10 Token Ring network is three times faster than existing ProNet-10 adapters, according to Proteon. The card connects IBM XTs and ATs to Proteon's 10-Mbps Token Ring network over shielded and unshielded twisted-pair (STP, UTP) cabling. Its interface is an 8-/16-bit adapter that supports most popular network operating software, including Novell NetWare 2.15, Banyan's VINES, and FTP Software's PC/TCP.

ProNet-10 supports UTP networking through an integrated media filter on the p1308 network adapter, an external UTP filter for the existing ProNet-10 installed base and for Proteon's Series 70 Intelligent Wire Centre.

The p1347, p1346, and p1342 adapter cards are high-performance bus master adapters that feature on-board media filters for UTP connections. The p1347 is an IBM AT file server adapter, the p1346 is an IBM AT workstation adapter, and the p1342 is an IBM PC and XT adapter.

Proteon's Series 70 Intelligent Wire Centre is a modular system available in mix-and-match configurations that supports 4-, 10-, and 16-Mbps networks. When you combine it with TokenView network management software, it provides capabilities for network management and configuration control. Two models are available: the p7201 for ProNet-4 and -16 networks over STP, and the p7202 for ProNet-4, -10, and -16 networks over UTP.

The Series 70 uses an independent, out-of-band communications channel for management command and control.
Price: £459 for the p1308; £495 for the p1347; £425 for the p1346; £385 for the p1342; £995 each for the p7201 and p7202.
Contact: Proteon International Ltd., The Business Park, Technology Dr., Beeston, Nottingham NG9 2ND, UK, 44-602-431311.
Inquiry 932.

Software for Sawmills

Ciris says that you can maximize the profitability of your sawmill with its Delta Series of software modules. The flexible modular design lets you choose and add modules to meet your company's changing needs, while the system's design flexibility lets you set production costs and output values for an individual mill.

The modules let you calculate the optimum (profit-maximizing) output of various board types from any given log, compute manufacturing costs and profit margins for any given production level, and control production output efficiently.

The Delta modules run on the IBM PC with 640K bytes of RAM, DOS 2.0, and a 20-MB hard disk drive. A LAN version is also available.
Price: 50,000 French francs and up.
Contact: Ciris, rue Thomas Edison, Z.I., 33600 Pessac, France, 33-56-07-22-20.
Inquiry 904.

continued

DHPaint lets you grab an image and then modify or retouch it in virtually any way conceivable, while displaying the effect as the work is done.

Video Painting Software for MicroEye TC

D igithurst's DHPaint video paint software is designed for use with the company's transputer-based MicroEye TC image grabbing and processing board for the IBM AT (see August 1989 What's New International, page 80IS-16).

You operate the package from a graphical icon menu, using a mouse, digitizing tablet, or other pointing device. It lets you grab an image and then modify or retouch it virtually any way you want, while displaying the effect as the work is done.

DHPaint gives you a full complement of freehand, line, Bézier curve, polygon, rectangle, circle, and text drawing tools with 32 preset brush shapes and the ability to modify a brush to any 32- by 32-pixel pattern. The software works with images of 720 by 512 pixels and 16.7 million colors and supports VGA display formats. You can choose colors from the screen, from a palette, or mixed from hue, saturation, and intensity or RGB values. There are 11 logic modes available when you are drawing, including tint and blend, that help you achieve a variety of effects and textures.

You can selectively lighten, darken, or intensify colors in a grabbed image and even completely rearrange the colors, producing psychedelic or pseudo-color effects on a moving image. You can also achieve pattern drawing, spraying, or filling by selecting a pattern area from the screen image. When you generate a pattern from an area of a grabbed photographic image, you can create realistic retouching effects. You can load and save images in TIFF

and TGA formats for compatibility with other applications, such as slide making and electronic publishing software.

A cut-out function lets you pick and save an arbitrarily shaped area of the image and then reposition, resize, and redraw it on the same or another image. This feature would be useful for applications such as photofitting and hairstyling in which you would call up a library of standard images from disk and try them in position until you achieved the desired result.

DHPaint's sophisticated stencil, or masking, function lets you achieve a variety of image effects by combining two images that are held in the two picture stores on the MicroEye TC board. For example, you can store an x-ray photograph of a hand in the reference store and capture a real image of a hand into the first store. You then cut an outline stencil in the image of the real hand and fill it with the same portion of the image of the x-rayed hand to achieve the effect of a cutaway showing the bone structure below the flesh. Alternatively, you can spray the x-ray image onto

the real hand image to achieve a ghostly effect of the bones showing through the hand.

DHPaint comes bundled with the MicroEye TC board. It requires 640K bytes of RAM and DOS 2.1.
Price: £3950.
Contact: Digithurst Ltd., Newark Close, Royston, Hertfordshire SG8 5HL, UK, 44-763-242955.
Inquiry 948.

High-Capacity Hard Disk Drives

I CM and CTS offer hard disk drives with capacities exceeding 100 MB. Both devices come with interfaces that are compatible with the NEC PC-9801-55 board. The ICM SRC-131 has a capacity of 130 MB, while the CTS HDa-PC110-55 has a capacity of 114 MB.
Price: 228,000 yen for the ICM SRC-131; 245,000 yen for the CTS HDa-PC110-55.
Contact: ICM, 4-2-9 Nihonbashi, Naniwa-ku, Osaka 556, Japan, 81-06-648-4702; CTS, Dai-ichi Kobayashi Building, 32-9 Taito, Taitoku, Tokyo 11, Japan, 81-03-835-1887.
Inquiry 951.

Image Analysis
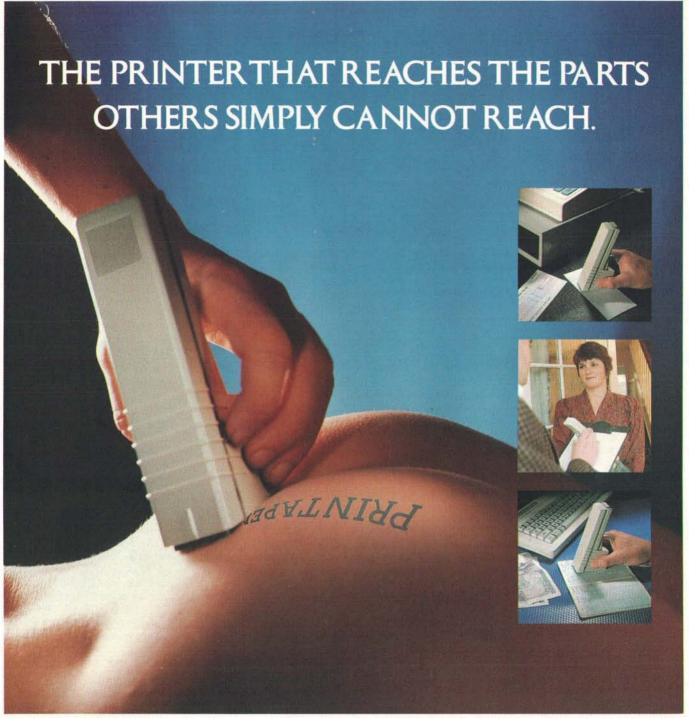
O FA Master is a personal-computer-based image-analysis system that you interface to factory or laboratory equipment through an integrated I/O board on the computer bus. You can use the system for industrial applications such as process monitoring and control or for laboratory applications such as laser-beam measurement and DNA sequencing.

The package includes a scientific-grade charge-coupled-device camera, a framestore, a 386 computer, and a 940-MB WORM (write once, read many times) disk drive, plus software applications and control packages.

The advanced framestore and digitizer can perform processing operations in hardware. You use the optical disk drive for image archiving, and it is integrated with a database for efficient retrieval of stored data. The BASIC-like programming language, developed by Oxford Framestore Applications, lets you store macro sequences as library routines for execution by the control program and program I/O control signals using the same syntax as the image-processing commands. Software application modules are available for automated image archiving, particle sizing, low-light-level imaging, object tracking, DNA sequencing, instrument readout, and laser-beam analysis.
Price: About £17,000.
Contact: Oxford Framestore Applications Ltd., Orchard End, Ickleton Rd., Wantage, Oxfordshire OX12 9JA, UK, 44-2357-66078.
Inquiry 905.

A Voice-Mail System for 50 to 400 Users

CatsCo's Telmail 1000 units are complete hardware and software voice messaging systems designed for small to medium-size businesses. The systems can operate stand-alone or fully integrated from a PABX. They provide voice-mail facilities such as telephone answering, message recording, replay, forwarding, and distribution.

The Telmail desktop processing unit (measuring 60 by 40 by 15 cm) contains a main Unix processing card, a voice-to-data interface card, a hard disk drive, and, if required, a telecommunications four-wire to two-wire extension converter card. Four models are available: the Model 1050 for 50 users and the Model 1100 for 100 users (both models with two voice channels), the Model 1200 for 200 users with four voice channels, and the Model 1400 for 400 users with eight voice channels. All models are compatible, so you can upgrade to a larger system at any time with little disruption to voice-mail services, according to the company.

The system can answer the phone with a personal greeting, record a message, or transfer the caller to a receptionist. Message replay features let you browse or replay messages in full, skip forward or backward within messages, adjust replay volume, erase messages, reply with comments, or forward messages on with comments. Message recording features include pause during record; review and rerecord; message identification stamps; confirmation of message delivery; and distribution to mailing lists. You can also connect the unit to terminals or

The four models in the Telmail 1000 range of voice messaging systems are available for from two to eight voice channels with 50 to 400 mailbox users.

printers to run off reports.
Price: £9000 for the Model 1050; £12,500 for the Model 1100; £14,500 for the Model 1200; £17,000 for the Model 1400.
Contact: Cats Company (CatsCo) (UK) Ltd., Abberley House, Great Shelford, Cambridge CB2 5LQ, UK, 44-223-841000.
Inquiry 920.

Add up to 8 MB of Memory to Your IBM PC

The S-RAMCard adds up to 8 MB of memory to your IBM PC, XT, AT, or PS/2 Model 30 with a maximum clock speed of 30 MHz. The full-size card offers 64 register sets, which let it swap applications in a multitasking environment faster and allow the card to receive data at 9600 bps in background mode, without dropping characters, while running another application in the foreground. In addition, the S-RAMCard provides hardware and software support for EMS 4.0 and EEMS.

Even before you install the board, the setup program lets you set S-RAMCard configuration parameters and then gives you a graphical repre-

sentation of the switch settings you will need. The setup program also installs the EMS driver.

The basic S-RAMCard comes with 2 MB of memory. You can upgrade it in 2-MB increments using 1-megabit DRAMs to achieve a maximum of 8 MB.
Price: 1995 Swiss francs for the basic 2-MB card; 2480 SF with 4 MB; 4535 SF with 8 MB.
Contact: Schnellmann Interhandels AG, Churerstrasse 160a, 8088 Pfaeffikon, Switzerland, 41-55-48-51-61.
Inquiry 930.

Add-Ins for Lotus 1-2-3

If you work in the options/investment field, you may find Option Price Calculations from 4-5-6 World helpful. The new add-in for Lotus 1-2-3 and Symphony helps you perform specific analytical tasks within the specialized area of options for many types of assets, stocks, bonds, currency, commodities, and futures contracts.

With Option Price Calculations, you can calculate the theoretical value of a stock option, the hedge ratio of a stock option, or the rate at which an option's price changes as time passes. You can even calculate the sensitiv-

ity of an option's price to changes in volatility, strike price, or holding costs.

The software gives options @functions that you can use like standard Lotus @functions; once they are attached, you can use them in any combination with Lotus @functions.

Option Price Calculations requires 44K bytes of RAM in an IBM PC and Lotus 1-2-3 release 2.01 or any release of Symphony.

Also available from 4-5-6 World is a new version of See-More, another of the company's add-ins for Lotus 1-2-3. SeeMore 2 lets you see more rows and columns, create multiple windows, and display live graphs in Lotus 1-2-3 release 2 or higher. The software gives you a choice of compressed formats that double, triple, or quadruple the number of cells you can display on-screen. In addition, you can zoom in or out with just a single keystroke.

Version 2 is a flexible windowing system that supports titled and overlapping windows, including multiple graph windows; the windows remain active while you operate Lotus 1-2-3. You can edit a macro in one window and monitor its operation in another. Graph windows are hot-linked to spreadsheet windows, so when you change a number in any spreadsheet window, all dependent windows are automatically updated.

SeeMore 2 runs on the IBM PC with 64K bytes of RAM over the requirement for Lotus 1-2-3 release 2.
Price: £395 for Option Price Calculations; £79 for SeeMore 2 (£35 as an upgrade).
Contact: 4-5-6 World, Wellington House, Butt Rd., Colchester, Essex CO3 3DA, UK, 44-206-579776.
Inquiry 917.

continued

Softline INTERNATIONAL DIVISION

This Leo Roars at 33 MHz

Centerprise International has added a 33-MHz system to its Leo range of personal computers. The new model is available in desktop, compact desktop, and floor-standing tower configurations.

The Leo 33-MHz 386 is based on a Micronics motherboard and can be run at either 33 MHz or 8 MHz to ensure software and hardware compatibility. It differs from the other Leo systems in that it is designed around a 32-bit bus with true 32-bit memory to produce a rating of 7.779 MIPS. The system performance is further enhanced by fast-access hard disk drives,

The Leo 386 is a 33-MHz computer available in floor-standing tower, desktop, and compact desktop configurations.

cache memory, and intelligent video controller cards.

The desktop configuration is designed for CAD/CAM and desktop publishing applications. It includes 2 MB of memory with direct-mapping

cache memory, DR-DOS 3.4, a VGA color display, the Leo VGA controller card, one 5¼-inch floppy disk drive, and a 40-MB hard disk drive.

The floor-standing tower configuration is designed to

act as a file server in a Novell networking environment or as a Unix or Xenix central controller. This model comes with 8 MB of cached RAM and a 160-MB hard disk drive.
Price: £3650.
Contact: Centerprise International Ltd., Hampshire International Business Park, Crockford Lane, Chineham, Basingstoke, Hampshire RG24 0GQ, UK, 44-256-463754.
Inquiry 912.

Speed Up Your Zenith Z-248

With Z-Master 386, you can replace the 286 microprocessor in your Zenith Z-248 computer with a
continued

true 32-bit 386 microprocessor operating at 20 MHz with zero wait states.

To install Z-Master 386, you simply lift out the existing 286 CPU card and insert Z-Master into the same slot (the upgrade will also fit the Zenith Z-286 and Z-241). The 386 upgrade employs Phoenix Technologies' ROM BIOS and supports operating environments like Microsoft Windows/386, OS/2, VP/ix, Unix, and Xenix 386. The power of your upgraded machine provides an ideal platform for software development, CAD/CAE applications, desktop publishing, spreadsheets, and office automation, according to the company.

You can incorporate up to 4 MB of 32-bit memory with a highly efficient cache circuitry allowing access not only to the on-board 32-bit memory, but also to all the existing 16-bit memory.

Z-Master comes with Disk Quick, a caching program that makes programs run faster by substantially reducing the number of physical disk-read operations. Other features include dual clocks for greater system bus compatibility and support for an Intel 80387 numeric coprocessor, providing floating-point performance of up to 10 times that of a Zenith Z-248 with an 80287 coprocessor.
Price: £1295 for a 1-MB version; £1795 for a 4-MB version; £398 for a 20-MHz 80387 coprocessor.
Contact: Corporate Software Ltd., Corporate House, 23 Horseshoe Park, Pangbourne, Reading, Berkshire RG8 7JW, UK, 44-734-845361
Inquiry 946.

Hard Disk Drives for the Mac

Formac offers a series of hard disk drives for the Macintosh with storage capacities ranging from 30 to 582 MB. Two of the models have disk caching.

The ProDisc 40DC and ProDisc 80DC give you storage capacities of 40 and 80 MB, respectively, an average access time of 19 ms, a data transfer rate of 16 Mbps, and 64K-byte disk caches. The drives automatically adapt themselves, without reformatting, to work at optimal speed on all Macs. The software lets you create up to 31 partitions for organizing your data and features password and write protection; DiscCloning to provide continuous, permanent backup; and AccelAdapt for optimizing performance with accelerator cards. The true partitioning lets you run a different operating system (Mac OS or Unix) on each partition. Internal versions of the ProDisc 40DC are available for the Mac II, SE, and Plus.

The ProDisc 30, 180, 300, and 600 hard disk drives give you storage capacities of 30, 172, 288, and 582 MB, respectively. The average access times of the drives range from 38 to 16 ms, and the data transfer rates range from 10 to 12 Mbps. The ProDisc 180, 300, and 600 drives come with the same software as the ProDisc 40DC and 80DC.
Price: 1800 and 2100 deutsche marks for the ProDisc 40DC; 3050 and 3350 DM for the ProDisc 80DC; 1290 to 10,900 DM for the ProDisc 30, 180, 300, and 600 drives.
Contact: Formac GmbH, Charlottenstrasse 13, D-1000 Berlin 61, West Germany, 49-30-251-04-01.
Inquiry 928.

CS 3616s

VEGAS

VEGAS

VEGAS COMPUTER COMMUNICATIONS GROUP

| London | Milan | Barcellona | Paris | Bonn | Athens | Moscow |

Being personal is good enough. Or is a personalized computer better?

A personalized computer is better.

Vegas computers are personalized. They can grow with the users' requirements since the only upgrade necessary is a CPU module when more speed or power is required for the users' applications.

Vegas computers are personalized because they can be supplied to match complex and unusual configurations on request.

Vegas computers are personalized because they can speak English, Italian, German, French, Spanish, Greek and Russian.

CS Family   CE Family   CF Family   CT Family

VEGAS

Adds wings to your thoughts.

Vegas Computer Communications Group references on following page.
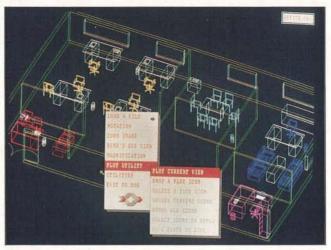
New Products for AutoCAD Users

Artist Graphics Europe offers two products for AutoCAD—Artist GT ReViewer, a software management tool set for AutoCAD designers, and Artist Dual GT, a high-performance driver that works with the Artist XJ10 high-resolution graphics controller to offer dual-screen workstation capabilities for CAD/CAM/CAE designers working with AutoCAD on the IBM AT.

The Artist GT ReViewer software consists of seven utilities in a drop-down menu format that let you review and annotate drawings in Auto-CAD files, even on machines not equipped with AutoCAD. It is designed for use with the IBM AT and compatibles equipped with the company's Artist XJ8, XJ10, and XJ12 graphics controllers as well as with the Artist 10MC and 12MC graphics controllers for the IBM PS/2s.

The Load A File utility lets you load GT Slide files or shaded files, and the Zoom Image utility lets you review those files using the pan or zoom functions. The Notation utility lets you annotate drawings at any zoom level—you place comments in a text window that closes to a bright yellow triangle to flag the notation.

You can view your entire drawing with the Bird's-Eye-View utility. This is especially useful when you are in zoom mode because the utility highlights the portion of the drawing that you have currently displayed on your screen. You can change the views instantly, as well as re-size and reposition them as required. The Bird's-Eye-View is automatically updated as you make changes and additions to the files.

The Artist GT ReViewer software consists of seven utilities that let you review and annotate drawings in AutoCAD files, even on machines not equipped with AutoCAD.

The Magnification feature lets you magnify a small section of the drawing—for instance, only the portion you are working on at the time.

With the Utilities function, you can create a file type that you can read with View-Base, a database for Auto-CAD. You can review existing files or scan additional drawings to create new files. This function also improves the display resolution over that produced by ViewBase alone, according to the company.

You can transmit files to an Artist Device Interface–compatible plotter almost instantaneously with the Plot utility, emulating a print screen function. This high-performance driver significantly reduces the time required to plot even the most complex drawings, thus freeing your system to perform other tasks.

In addition, the speed of the Plot utility makes checking plots a practical alternative if you want to review, quickly

edit, or annotate hard-copy AutoCAD drawings. You can also check the entire drawing or only a portion of it in draft quality.

The Artist GT Plot driver supports AutoDesk's Fast Plot interface and is available separately for use with systems equipped with the Artist GT XJ10 graphics controller.

The Artist Dual GT driver is designed to enhance the capabilities of systems currently configured with dual screens and EGA or VGA graphics cards. You use one full-size screen for creating and editing and a second low-resolution VGA monitor to provide a continuously updated Bird's-Eye-View for panning and zooming; this screen also displays text. You can use the Bird's-Eye-View on the lower-resolution monitor as a reference for maintaining location as you work in greater detail on the high-resolution monitor.
Price: £295 for Artist GT ReViewer; £70 for the GT Plot Driver; £250 for the Artist Dual GT driver.
Contact: Artist Graphics Europe, B.I.M. House, Crofton Rd., Dun Laoghaire, Dublin, Ireland, 353-1-800919.
Inquiry 914.

Tektronix's Entry-Level Graphics Superworkstation

Tektronix offers an entry-level addition to its XD88 range of RISC-based graphics superworkstations. The XD88/10 is a high-resolution bit-mapped workstation that can operate at a processing speed of 17 MIPS. The company designed the workstation for earth resources and geographical information systems, mechanical CAE, manufacturing automation and process control, and data analysis.

The XD88 family incorporates two- and three-dimensional workstations, an application processor, and a server node. Based on Motorola's 88000 RISC chip, all the systems have Ethernet LAN capability, two RS-232C serial ports, a Centronics parallel port, and a high-speed SCSI port. The XD88/10 also comes with a 16-inch color monitor (a 19-inch color monitor is an option), a 156-MB hard drive (expandable to 3 gigabytes), and 8 MB of RAM (expandable to 32 MB). It is compatible with Tektronix's XD88/20 and XD88/30 workstations. The workstation's 256-color system is expandable to 24 bit planes with an 8-bit overlay. In addition, you get a VME slot for special functions.

The XD88/10 utilizes a Tektronix proprietary gate-array CMOS IC, called Color=Cache, that optimizes color bit-mapped performance by storing the colors you are currently using. Color=Cache chips provide an interface between the workstation's 32-bit system bus and the 256-bit Pixel Data Bus that goes to the frame buffer (the area of memory that holds pixel-by-pixel display information). Once you store

continued

VEGAS
COMPUTER COMMUNICATIONS
GROUP

UNITED KINGDOM
Unit 2 - 243, Kangley Bridge Road
SE26 5BA Sydenham - London
Tel. 44-1- 6599166
Fax 44-1-7782469

ITALY
50, Via Casignolo - 20092 Cinisello Balsamo (Milano)
Tel. 39-2-66010306/66010406
Fax 39-2-66010598

149, Via Baldanzese - 50041 Calenzano (Firenze)
Tel. 39-55-8878151
Fax 39-55-8878159

SPAIN
210, 1̃, 6̃, Calle Aragon - 08011 Barcellona
Tel. 34-3-4513383/93
Fax 34-3-2538163

FRANCE
38, R. des Estats Generaux - 78000 Versailles
Tel. 33-1-39511011
Fax 33-1-39532030

WEST GERMANY
Siegenerstrasse 23 - 5230 Altenkirchen
Tel. 49-2681-3041/2
Fax 49-2681-7222

GREECE
35, Ippokratos St. - 167 77 Elinikon, Athens
Tel. 30-1-9910950
Fax 30-1-9932801

SOVIET UNION
36, Obrouchev St. - 117342 Moscow
Tlx 412160 ANTAR-SU
Fax 7-95-4202250

✘VEGAS

SOFTWARE SECURITY MODULE

An easy to use
hardware/software
combination for securing
data/programs.

Our security modules
can be used simply as
'Dongles' through to
advanced encryption –
decryption systems.

★ For IBM PC family (and compatibles)
★ Variants for other systems
★ Advanced construction – rugged and compact
★ Proven throughout the world
★ Transparent operation – user friendly
★ Driver software supplied for variety of applications

master charge

VISA

CTL Control Telemetry of London
11 Canfield Place, London NW6 3BT
Tel: 01-328 1155 Fax: +441 328 9149

USE YOUR PC'S MORE EFFECTIVELY!

- **Extend the distance between Keyboard/Monitor and CPU up to 600 feet with:**
  EXTENDER

- **Connect 2 or 3 Keyboards and Monitors to 1 CPU with:**
  COMPANION

- **Control up to 8 PC File Servers with a single Keyboard and Monitor with:**
  COMMANDER

- **Boot up your PC by phone with:**
  PHONEBOOT

- **Reboot your PC File Servers by phone with:**
  REBOOT

Cybex Corporation

2800H Bob Wallace Ave. • Huntsville, AL 35805
(205) 534-0011 • FAX (205) 534-0010

---

the colors, the processor can continue drawing with them by sending only a single bit of data, rather than the full 8 bits that would ordinarily be needed to specify one of the 256 displayable colors. According to the company, this results in a data transfer rate that is up to eight times faster than that of systems without the Color=Cache IC.

With the transparency function, the processor can select only the pixels that need to be modified. The frame-buffer scrolling model produces high-speed windowing and quickly moves graphics from one area of the screen to another. In addition, a unique Color=Cache mode allows the CPU to treat the frame buffer as if it contained a single bit per pixel, so monochrome applications can run without modification or performance penalty.

The XD88/10 supports Tektronix's high-speed Color=Cache implementation of the X Window System 11.3—it runs software ported to X11 on Tektronix's other XD88 superworkstations or its 4319 workstation, with little or no modification, plus software that is compatible with 88000 BCD standards.

The workstation also supports UTek V, Tektronix's implementation of Unix based on System V with BSD 4.3 extensions and SVID and POSIX compliance, Ethernet with TCP/IP, binary compatibility standard, and PC emulation. The PC emulation allows DOS applications to run on a window in the workstation. Price: £14,396.
Contact: Tektronix UK Ltd., Fourth Ave., Globe Park, Marlow, Buckinghamshire SL7 1YD, UK, 44-6284-6000.
Inquiry 922.

Transtech Devices' Graphics TRAM

Now you can get the Inmos Colour Video Controller, the IMS G300, in a transputer module format. Transtech Devices' TTG3 Graphics TRAM lets you integrate a flexible and programmable graphics subsystem into a transputer-based motherboard.

The TRAM consists of a 25-MHz T800 transputer that is linked to three functional blocks by a simple bus. The blocks include the G300 Colour Video Controller, 2 MB of video RAM (VRAM), which can be double-buffered where you need alternate screen banks, and 2 MB of DRAM (expandable using 4-Mb devices) for program memory and screen manipulation workspace. You can integrate the TRAM into parallel systems based on IBM PCs to Sun workstations.

The TTG3 supports screen resolutions of up to 1280 by 1024 8-bit pixels with a palette of 256 colors from a choice of 16.7 million shades; the RGB plus Sync outputs are ready for connection to standard and high-performance color monitors.

With its video timing generator, VRAM interface, color lookup table, and 3- by 8-bit D/A converters, the G300 is a fully programmable video controller that gives outputs compatible with RS-170a and EIA-343. The package is supported by a set of Occam primitives that generate key graphics functions, which you call from the compilers or the development system running the application. This allows the TTG3 to either implement graphics primitives directly or act as an intelligent channel receiving data via the bidirectional transputer links, at either 10 or 20 MBps.

Price: £3295 and up.
Contact: Transtech Devices Ltd., Unit 17, Wye Industrial Estate, London Rd., High Wycombe, Buckinghamshire HP11 1LH, UK, 44-494-464303.
Inquiry 941.

Develop Rule-Based Expert Systems

You can use Nemo 2.0 to monitor an ongoing process and respond to events as they occur through its Perception Module, which handles all kinds of data acquisition from supervisors, databases (e.g., Oracle and Ingres), programmable logic controllers, and sensors. You can use the package for critical maintenance functions and safety evaluation, allowing corrective and predictive diagnosis in fields such as the aerospace industry, and where the volume and complexity of information make it impossible for an operator to make correct timely decisions.

Written in C, Nemo 2.0 is a rule-based expert-system development package for factory-floor applications such as alarm filtering, adaptive control, and fault diagnosis or advisory. It consists of a facts base, a rule base, and an inference engine, which employs techniques like temporal logic, forward chaining, and nonmonotone reasoning. The knowledge format includes rules and frame-like representation, as well as object hierarchy trees—production rules accept internal actions and priorities and can generate external actions.

The product development environment includes a Knowledge Compiler and Optimizer, a graphical multiwindow Developer Interface, and a graphical multiwindow Operator Interface. You can develop both interfaces for ASCII displays or high-resolution color graphics displays. In addition, they support SunTools and version 11.3 of the X Window System.

The compiler detects errors in the knowledge description and compiles the knowledge in such a way that Nemo's inference engine can make optimum use of it. The Developer Interface lets you build an application efficiently, using a multiwindow system operating on a bit-mapped or alphanumeric screen. It controls the inference engine, knowledge inspection, explanation of facts, knowledge tree, fact transmission, and answers to questions.

You can develop the Operator Interface by means of A-Views, a system for constructing interfaces for display and transmission of facts. You construct the screens with a file describing the different views (e.g., title and position) and fields (e.g., label and value type) associated with each view. A help menu is automatically associated with each value.

The Perception Module acquires signals from the process, smooths noisy signals, performs symbolic processing of measurements (e.g., thresholds, trends, and alarms), performs processes specific to the reasoning mode, and sends facts to the inference engine.

Nemo 2.0 runs on a 386 workstation under Unix; on the Sun-3, Sun-4, and SPARC-Station 1; on the Hewlett-Packard 9000 Series; and on the DEC VAXstation with Ultrix or VMS. The minimal configuration required is 1 MB of RAM and 3 MB of available space on a hard disk drive.
Price: 100,000 French francs.
Contact: S₂O Développement, 14, rue du Théâtre, 75015 Paris, France, 33-1-40-59-44-10.
Inquiry 949.

continued

NOW YOU SEE THEM...

BATCH SERVER

PRINT SERVER

NOVELL FILE SERVER

E-MAIL SERVER

DIAL-IN SERVER

DIAL-OUT SERVER

NOW YOU DON'T!

WITHIN THE NOVELL FILESERVER....
- BATCH SERVER
- PRINT SERVER
- E-MAIL SERVER
- DIAL-IN SERVER
- DIAL-OUT SERVER

NOVELL FILE SERVER

DEDICATED SERVER SYSTEMS

Dedicated server systems are a necessary part of today's high-powered LAN. Dedicated servers distribute the power of individual processors to specific LAN tasks. They protect your fileserver from the pitfalls of unexpected crashes, and increase performance by offloading tasks from your main system processor.

A GOOD IDEA—*But look what you'll need:*

- A standalone computer for each dedicated function.
- Extra networking hardware and expensive cabling to connect them together.
- A spare piece of office real estate in which to keep them.
- An explanation why you can't use your expensive computers for anything else.

DEDICATED SERVER BOARDS

QL Series boards—one, two, or four XT- and one or two AT-compatible computers on a card—are elegant replacements for dedicated, standalone machines. They have all the performance and processing power of a computer, but plug into the bus of your fileserver, bridged system, or expansion chassis—out of sight, ready to process tasks on demand.

A BETTER IDEA—*Look what you can do:*

- **ASYNCH DIAL-IN**—With QL boards and communications software, you can dial into network resources from your home or remote office PC.
- **ASYNCH DIAL-OUT**—Dial out using communications software from any network node to bulletin boards, information services, and other networks.
- **DEDICATED FILE SERVICES**—Manage unattended tasks such as message handling for electronic mail, background batch processing, and print queuing.

IBM, PC/AT, and PC/XT are registered trademarks of International Business Machines Corporation

For better dedicated server solutions, call Cubix Corporation today at 1-800-829-0550

≡CUBIX
CORPORATION

Cubix Corporate Offices • 2800 Lockheed Way, Carson City, Nevada 89706 U.S.A. • Tel (702) 883-7611 • Fax (702) 882-2407
Europe • Unit 4 Colonial Business Park, Watford, Hertfordshire, WD2 4PR, England • Tel (44) 923 51150 • Fax (44) 923 37021

Circle 417 on Reader Service Card (DEALERS: 418)

Paradise for Transputer Storage Problems
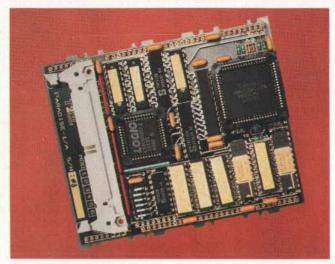
T2 Systems' Paradise series of transputer/SCSI systems can link to disk drives from a 10-MB hard disk drive up to the largest WORM (write once, read many times) drive. You can also link the systems to tape streamers and SCSI-based printers and plotters.

The Paradise-1 SCSI transputer module (TRAM) is a dedicated 10-server system based on the T222 transputer coupled to 16K bytes of static RAM and 8K bytes of fast bipolar PROM. The SCSI port section uses a 53C80 SCSI controller connected through removable terminators to a fully latched low-profile 50-way connector. Also, by making use of the transputer's own concurrency, you can transfer a data block of any size with a single command at a sustained rate of 700K bytes per second.

When you install T2 Systems' I-connect software on the Paradise-1, you can easily interface to SCSI peripherals. You fill up a control block of five integers, with appropriate values gleaned from the peripheral's manual, and call an interface routine called I-talk. The I-connect software then arbitrates for the SCSI bus, selects the peripheral, and executes the command, filling or emptying a byte-wide data buffer as necessary.

I-talk is available for a variety of languages, including Occam and the Inmos and 3L C compilers.
Price: £795 for Paradise-1; £195 for I-connect.
Contact: T2 Systems Ltd.,

The Paradise-1 SCSI TRAM is a dedicated 10-server system based on the T222 transputer coupled to 16K bytes of static RAM and 8K bytes of fast bipolar PROM.

62 Longmead Ave., Bishopstoke, Eastleigh, Hampshire SO5 6ET, UK, 44-703-641276.
Inquiry 921.

Icon-Oriented Computer Access for Windows

PubTech File Organiser is the latest addition to Neow's range of Microsoft Windows software. The icon-oriented computer access system, which runs under Microsoft Windows/286 and Windows/386, offers more than 80 file-manipulation functions.

The software lets you manage files and folders (directories) via its visual iconic representation. By simply clicking and dragging icons across your screen, you can create, rename, move, and copy files and directories. In addition, you can delete a file or group of files simply by dragging them to the garbage icon. To print, you simply drag files to the printer icon.

PubTech File Organiser's utilities help you manage virtually all aspects of your computer environment. For example, you can recover files that you have deleted under any

software package and easily find files or text within files anywhere on your disk, even if you've forgotten their names or location.

The ASCII editor contains the typical features of an editor, such as Open, Copy, Cut, Paste, Save, and Revert, and it can automatically save files at specified time intervals while you are editing. You can remap your keyboard to emulate several popular software packages or choose your own key combinations. You also have access to a macro language, an appendable clipboard, and an unlimited Undo function.

You can customize your desktop so that the files and applications you use most are instantly available when you start working on your computer. When you need to run different applications on your screen simultaneously, the software lets you control the placement of the applications. Any windows or applications that you leave open when you leave PubTech File Organiser can be optionally saved and reopened when you return to PubTech.

You select disk functions such as Diskcopy, Make System Disk, Format Data Disk, and Erase Disk from pull-down menus. The functions support a variety of disk drives and disk types, including drives attached to networks. The package also provides a complete backup and restore function.

PubTech File Organiser requires an IBM AT with 512K bytes of RAM, DOS 3.0, Microsoft Windows 2.0, a mouse, and a hard disk drive.
Price: £145.
Contact: Neow, Unit 12, Progress Business Centre, Whittle Pkwy., Slough SL1 6DQ, UK, 44-6286-68334.
Inquiry 942.

Oki Electric's 25-MHz 386 Machine

Oki Electric's if-800 EXII 110/120 25-MHz 386 machines have the same basic architecture as the if-800 EX, with 4 to 20 MB of main memory and a 64K-byte cache. As with its EX machine, the company considered overall system balance and used multiplexed buses for the system bus, memory bus, and I/O bus. The designers placed the frame memory between the CRT display hardware and video RAM and used cell mapping to boost display speed.

At the same time, Oki Electric announced Microsoft Windows 2.1 and the db-1HS upgraded database server. The company is changing the CPU to the 25-MHz 386 and providing a disk cache to boost performance by a factor of at least 10.
Price: 798,000 to 1,188,000 yen.
Contact: Oki Electric Industry Co. Ltd., 1-7-12 Toranomon, Minato-ku, Tokyo 105, Japan, 81-03-501-3111.
Inquiry 950.

Lucky-Goldstar, one of the 4 major conglomerates in Korea, proudly presents... **Goldstar**, a new brand with over 30 years of experience.

Goldstar has been a part of the electronics revolution, producing a complete spectrum of consumer and industrial products, often for major foreign OEMs.

Now we're introducing our own line of IBM compatible computers. From entry-level XT®s to more powerful desktop and laptop AT®s to 386™ file servers and PS/2® model 55's. And monitors in CGA, EGA, VGA, CAD/CAM, multi-scanning, and DTP models to match your customer's needs.

Goldstar expertise doesn't end at the computer screen. We make office automation equipment too. From fax machines to LBPs (lazer beam printers) to electronic typewriters. Products full of quality and value. All at competitive prices.

There's a new "star" on the office automation horizon. One that has been there. One that will be there in the future. Reliable, modern, innovative.

All names mentioned are trademarks or registered trademarks of their respective companies.
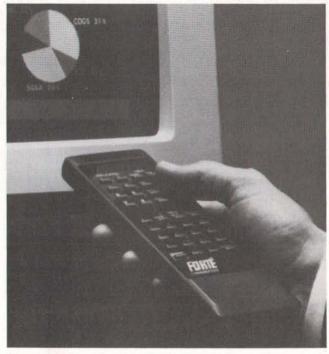
Control Your IBM PC from 50 Feet Away

Have you ever wished you had a remote keyboard for your computer to help you give a presentation or demonstration or to instruct students? If so, you should check out Azcom's Remote Keyboard, a hand-held remote keyboard for use with the IBM PC XT, AT, PS/2s, and compatibles.

Remote Keyboard operates by transmitting infrared signals to a receiver, which plugs into the computer's serial port. The software to run the remote is easy to install and is totally transparent to the application.

The device lets you control

Remote Keyboard is a hand-held remote keyboard for use with an IBM PC or compatible.

application programs from up to 50 feet away from your computer and gives you full keyboard functionality. It contains 40 user-programmable keys, plus additional keys for the PageUp, PageDown, Home, and reverse backslash functions. You can use the built-in software pointer for graphics presentations—you activate and control it with the touch of one button.

Remote Keyboard requires an IBM PC with 512K bytes of RAM, DOS 2.1, and one RS-232C serial port.
Price: £349.
Contact: Azcom Marketing Co., Selborne House, Old Ave., Weybridge, Surrey KT13 0PQ, UK, 44-932-851972.
Inquiry 918.

Matrix Layout 2.0

Fantasy or Reality?

Imagine a development system that takes your ideas and moulds them into applications that can include WIMPS, CASE, OOPS and HYPERTEXT.
And then has the ability to show you how the application works, looks and feels before it codes it up into Microsoft C, Turbo Pascal, Lattice C, Turbo C, QuickBasic or an .EXE file at the touch of a button.

Matrix Layout turning your ideas into reality

10.750 BF excl. VAT & delivery

MATRIX
EUROPE

Circle 445 on Reader Service Card

Matrix Europe N.V.   Geldenaaksebaan 476   B-3030 Leuven Belgium   Tel. +32 16 20 20 64   Fax +32 16 20 77 10

Please send me more information and a free demonstration disk

NAME	TITLE	
COMPANY		
STREET	POSTCODE	CITY
COUNTRY	TEL.	FAX

Parallel Processing for Sun Workstations

Transtech Devices' MCP1000 Multi-Computing Platform is a parallel processing system for Sun-3 and Sun-4 workstations that supports up to 32 transputer modules (TRAMs).

The transputer-based motherboard is structured into four areas: It connects a 32-bit VME interface into a processor farm consisting of four sites, each having eight TRAM slots. The processor farm feeds into a transputer link-switching matrix and then into a bank of link outputs, each giving direct I/O at data rates of 10 or 20 MBps. When fully loaded with 32 TRAMs, the MCP1000 is capable of 72 MFLOPS; with up to 12 platforms per Sun chassis, this offers a computing facility of 864 MFLOPS.

Up to four users can access each Multi-Computing Platform from any point in a Sun network; all members of the network could carry the MCP1000, thus giving a distributed computing resource running into gigaFLOPS. Fully integrated into SunOS Unix and NFS, the Multi-Computing Platforms are capable of working with a range of application-specific TRAMs, custom hardware, and standard peripherals.

The Multi-Computing Platform is supported by a parallel operating system called Genesys and a variety of Unix-based software tools; Genesys resides alongside SunOS.
Price: £4000 and up.
Contact: Transtech Devices Ltd., Unit 17, Wye Industrial Estate, London Rd., High Wycombe, Buckinghamshire HP11 1LH, UK, 44-494-464303.
Inquiry 936.

A motherboard for Sun-3 and Sun-4 workstations, the MCP1000 is a parallel processing system that supports up to 32 transputer modules.

New and Improved K-Spread for the Atari ST

Aimed at professional users, K-Spread 4 for the Atari ST offers more than 200 functions, including statistics, trigonometry, string, data management, error trapping, date and time, and logical and conditional functions. The program's macro facilities let you carry out repetitive operations, and the Record feature lets you experiment with macros.

In addition to spreadsheet functions, K-Spread 4 also has graphing and forms-design features to enhance the presentation and understanding of data. You can produce line graphs and horizontal, vertical, clustered, and stacked bar charts with hot-key operations, and the package updates graphs as you change data. In addition, you can load and save Lotus WKS and WKI, ASCII, and DIF files into K-Spread 4.

You can design forms by surrounding cells or blocks of cells with lines of varying thicknesses. These surround boxes, used in conjunction with the grid on/off and column/row header selector, can deliver good-quality output directly to a printer or, via GDOS metafiles, to desktop publishing and other presentation packages.

GDOS is included in the K-Spread 4 package, along with printer drivers suitable for many 9- and 24-pin dot-matrix printers; Atari laser printer drivers are available on request to K-Spread 4 owners. The GDOS drivers let you print sideways on many printers and give you improved print quality. By selecting smaller fonts, you can increase the number of cells displayed on the screen.
Price: £129.95.
Contact: Kuma Computers Ltd., 12 Horseshoe Park, Pangbourne, Berkshire RG8 7JW, UK, 44-734-844335.
Inquiry 915.
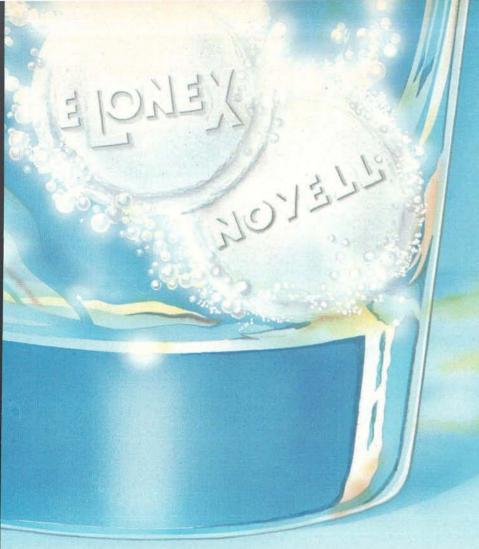
Hard Disk Drives with ESDI and SCSI Controllers

A new range of 5¼-inch hard disk drives is available from Phase IV Systems. The Micropolis 1560 and 1580 Series drives provide high storage capacity plus fast access and data transfer. The 765-MB drives can transfer data at 15 MHz with an average access time of 16 ms.

The 1560 Series features an ESDI controller, while the 1580 Series is equipped with a SCSI controller, which is designed to provide high-performance features such as synchronous data transfers at up to 4 MBps.

Using an advanced servo system, the 1560 and 1580 Series drives provide superior on-track and seek performance across a wide range of environmental conditions, according to the company. The dedicated disk surface for servo data leaves other surfaces free exclusively for user data, thereby eliminating format constraints and permitting more user data capacity as compared to embedded servo systems.
Price: £2700 and up.
Contact: Phase IV Systems, Unit 6, Oxford Business Centre, Osney Lane, Oxford, Oxfordshire OX1 1TB, UK, 44-865-792200.
Inquiry 935.

WHAT'S NEW

INTERNATIONAL

The Newest Version of InstaPlan

InstaPlan 5000 is the latest member of the InstaPlan family of project management software. It is file- and function-compatible with InstaPlan 2 (see September 1989 What's New International, page 80IS-62) and contains a Tracker option for recording actuals on your plans, a Notepad option for appending information to any activity or resource at any stage in your projects, and a Classic PERT option for printing PERT charts.

New features include single-keystroke consolidation with full retention of resource detail, no requirement for concurrent updating of source files, and full error reporting. The new version also supports EMS, taking the maximum plan size from around 500 activities to more than 10,000 activities.

InstaPlan 5000's Multifile Append facility can gather individual project plans and consolidate them for departmental or even company-wide review. It can collapse and expand groups of activities, making it easier to work with large plans because you can retain the detail on the area you are working on while viewing the rest of your plan in summary form.

The package also stores resource details, so you can review all the parts of a project and recognize and resolve staffing conflicts and project priorities before they become crises.

You apply the Resource Group Leveling feature to groups of resources to eliminate overloads, the Priority Leveling feature to multiproject plans, and the Localized Leveling feature to individual projects within a multiproject plan. With the Adjust Loading feature, you can delay resource commitments to low-priority projects, allowing high-priority projects to proceed. The Locate Overload facility lets you quickly scan future periods for a resource and identify any overloads.

With InstaPlan's Focus, Expansion, and Hide facilities, you can select activities and resource columns from the InstaPlan spreadsheet that you want to export in Lotus 1-2-3 or dBASE form to help you produce summary-level graphs and pie charts.

InstaPlan 5000 requires an IBM XT with 640K bytes of RAM, DOS 2.1, and 3 MB of hard disk space. The package supports up to 8 MB of EMS memory. It will run on networks using network shells like Novell NetWare, provided there is sufficient low-level DOS memory available to load InstaPlan (500K bytes) and the terminal has EMS memory.
Price: £499.
Contact: Deepak Sareen Associates, 28 Preston Rd., Wembley HA9 8JY, UK, 44-1-904-0811.
Inquiry 911.

continued

Tape Backup Systems for the IBM PC

The Retriever/60 and Retriever/60E are the latest models in Alloy Computer Products' Retriever range of tape backup systems. Both systems provide up to 60 MB of data storage on a quarter-inch cartridge, utilizing the QIC-40 standard error-correcting tape system.

The systems offer Error Correction Code, assuring data retrieval from old and worn tapes. Other features include file-oriented backup, DOS and Novell NetWare support, Alloy's own 386/MultiWare system support, and a scheduling utility for unattended single-user, multi-user, or network backup.

The Retriever/60 is an internal tape backup system designed for use with the IBM AT. It is compatible with most built-in floppy disk drive controllers and comes as a 5¼-inch half-height unit.

You can mount the Retriever/60E external tape drive horizontally or vertically. It

You can store up to 60 MB of data on a quarter-inch cartridge with the Retriever/60 and Retriever/60E tape backup systems.

derives its power from the host computer via a tape controller adapter board, which is compatible with the IBM PC, XT, AT, and PS/2 Model 30. The company also offers optional controller cards that are compatible with the other IBM PS/2 models.

The Retriever range includes the Tape System/2 internal backup system, which has been upgraded to 60 MB and is compatible with the IBM PS/2 Models 30, 50, 60, 70, and 80. The range also includes the Retriever/120, a 120-MB external tape system that is ideal for unattended

backup of LANs or multiuser system file servers.

All the Retriever tape drives include Alloy's ResQ and ResQNET software. ResQ lets you automatically back up and restore all files, disks, tagged files, or specified filenames and file groups; display a directory of files; erase or retain a tape; show vital tape statistics; or build an AUTOEXEC file. ResQNET allows backup on active Novell NetWare networks and preserves NetWare data security structures.
Price: £345 for the Retriever/60; £595 for the

Retriever/60E.
Contact: Alloy Computer Products (Europe) Ltd., Hurstleigh, Coronation Rd., Ascot, Berkshire SL5 9EY, UK, 44-990-872979.
Inquiry 913.

Speed Up Sun386i Workstations

Would you like to increase the performance of your Sun386i workstation from between 2½ to four times? Now you can with a new Intel i486 CPU upgrade board from Sun Microsystems.

The upgrade board features a 25-MHz CPU rated at 12 MIPS, an i486 chip, and enhanced I/O components. The upgrade entails no modification of your workstation other than swapping its 386 CPU board.
Price: £3950.
Contact: Sun Microsystems Ltd., Sun House, 31–41 Pembroke Broadway, Camberley, Surrey GU15 3XD, UK, 44-276-62111.
Inquiry 908.

If your software reads like an open book...

shut it.

HOW TO PUT YOUR PROFITS UNDER LOCK AND SOFTWARE KEY

Here's a simple fact.

Everytime someone pirates your software you lose money.

Every year software publishers lose hundreds of millions of pounds in this way. Then there are the problems of undetectable overseas piracy and the impossibility of taking legal action.

You can stop it happening to you by using the Software Key. As a micro-processor based, electronic device, it works by generating two random number sequences that don't repeat themselves for about 1.2 million years!

You can use these sequences to encrypt your programs, making it almost impossible for the pirate to pirate.

What's more, you can use the Software Key to encrypt data before disk storage, and before transmission down a dataline. It's also transparent, easily removable and portable.

If you're a bonafide software house and would like to see how the Software Key works, we can send you a sample pack for evaluation purposes.

Simply clip your business card or letterhead to the coupon stating your machine type and send it to:

Bristol Software Factory Ltd, 2 Southfield Road, Westbury-on-Trym, Bristol BS9 3BH, ENGLAND.

Or call Bristol (0272) 629790. Telex 94016310 (BSFL G) ELECTRONIC MAIL TGOLD 76:TWB010.

BY
BRISTOL
SOFTWARE
FACTORY
BSF

S·O·F·T·W·A·R·E
—— K·E·Y ——

A Window-Based Debugger for Parallel Programming

3 L designed Tbug, its window-based interactive debugger, to help parallel programmers observe and control the execution of their programs and thereby accelerate the debugging process.

The product's windowing environment lets you interact with concurrently executing tasks at the source level. You can freely run each task or dynamically created parallel execution thread in the system, single-stepped or halted at breakpoints. If you need access below the source level, Tbug also provides control down to the assembly language and register level.

You can examine memory values and change them using symbolic names or absolute addresses, and the software memory protection function catches tasks that attempt to overwrite code, modify another task's data, or access nonexistent memory.

Tbug requires an IBM PC with a B004 transputer and DOS 3.0. Although Tbug is primarily aimed at Parallel C 2.1, you can use it with previous versions of Parallel C, as well as with Parallel FORTRAN and Parallel Pascal.
Price: £220.
Contact: 3L Ltd., Peel House, Ladywell, Livingston EH54 6AG, UK, 44-506-41-59-59.
Inquiry 910.

Do-It-Yourself CD-ROMs

P ergamon Compact Solution offers CD Kreator, a CD-ROM toolkit that helps organizations construct their own CD-ROM publications for distribution in-house on

Tbug helps parallel programmers observe and control the execution of their programs and thereby accelerate the debugging process.

compact disks.

The kit supports a subset of the Standard Generalized Markup Language, which you use to place tags in the text to define data elements, such as titles, text, notes, bibliographies, references, citations, and document outlines. You then use CD Kreator's input filter software to process the SGML-coded text. The kit also provides indexing programs that support full-text searching, retrieval by specific data elements, and hypertext links.

The package includes two versions of the Knowledge Retrieval System, a CD-ROM search-and-retrieval software system: Text KRS handles text-only databases, and Graphic KRS handles text and graphics. You also get several tools that help you customize their user interfaces: RSC2 lets you modify the user interface and configure the resource files; KRS Preference File lets you establish and save preferences and database parameters; and GEMPaint, GEMDraw, and GEMScan let you create title screens.

CD Kreator runs on the

Macintosh with Finder 5.0 and on the IBM AT with 640K bytes of RAM, DOS 3.2, and 1 MB of EMS memory. Both versions require a 40-MB hard disk drive. The company also recommends a nine-track tape drive and a dedicated CD-ROM formatting and simulation system that includes disk storage and a tape drive.
Price: £9950.
Contact: Pergamon Compact Solution, Irwin House, 118 Southwark St., London SE1 0SW, UK, 44-1-928-1404.
Inquiry 902.

Portable User-Interface Development

W ith Interacter, you can create an interface with presentation graphics and point-and-shoot menus that is truly portable, according to Polyhedron Software. Interacter is a FORTRAN 77 user interface development system with some 250 text and graphics screen-handling routines.

Graphics features include a set of graphics primitives including input from a mouse or keyboard; Hershey fonts; Hewlett-Packard Graphics

Language plotter output/replay; presentation graphics routines for bar charts; three-dimensional function plots; a user-defined coordinate system; user-definable character sets and symbols; color, line style, palette, and plot-mode selection; and bit-block move and copy.

Text features include window management, point-and-shoot menus, and input boxes, all with pop-up features, even on serial terminals. You can switch applications freely between text and graphics modes. Other features include single-key input with device-independent keypad; cursor and function key detection; fixed-field string and numeric input with field editing; redefinable input control keys; character-level graphics for frames and diagrams in nongraphical screen modes; cursor control; attribute control; selective screen area clearing; and text and numeric field output.

The operating-system interface gives you directory lists, command-line execution, program chaining, operating-system variables, and command execution.
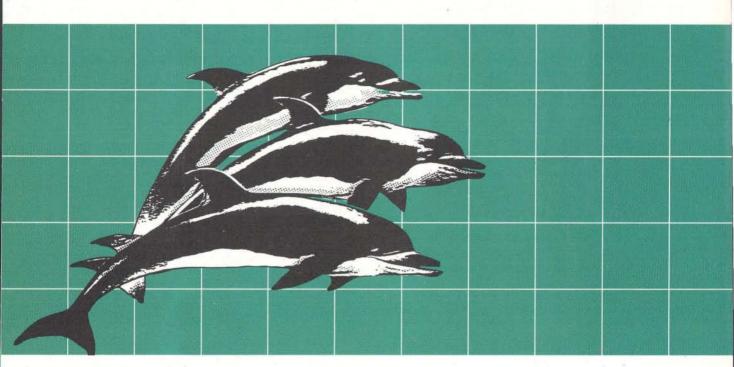
You can port software you have written with Interacter among Unix workstations, VAX or Prime minicomputers, IBM PCs, and the Acorn Archimedes, according to the company.
Price: £99 to £595 for the small systems version; £1000 to £4000 for the workstation and multiuser versions.
Contact: Polyhedron Software Ltd., Magdalen House, 98 Abingdon Rd., Standlake, Witney, Oxfordshire OX8 7RN, UK, 44-865-300-579.
Inquiry 931.

A Card for Digital Signal Processing

You can use the ATD-C25 for applications such as math accelerators, speech processing, image reconstruction and analysis, satellite communication, and radar and sonar analysis.

The ATD-C25 is a signal processing and program development system on a single card for the IBM PC with an 8- or 16-bit PC bus channel. The card includes an on-board 10-MIPS TMS320C25 digital signal processor (DSP) chip; a 256K-byte on-board data buffer (with zero-wait-state RAM), which you can address directly as program or data storage; a half-overflow detector for fast Fourier transforms (FFT); and 16 I/O expansion ports.

Included with the card is the DBG-C25, a debugger program for software development on the TMS320C25. The program supports single-step, breakpoint, and full-speed operation. It lets you examine registers and memory and modify them using a variety of formats. It also includes a disassembler. The commands of the DBG-C25 provide a real-time graphical display for showing DSP results. The program could also convert the chip's object code to Intel hexadecimal format for EPROM or PROM programming.

ATD-C25 supports the FFT spectrum-analysis package, called ATD-FFT, which you can interface to BASIC, FORTRAN, Pascal, and C.

Another software package, called the ATD-FID, consists of four interactive filter design modules capable of performing functions such as designing finite-impulse response and infinite-impulse response digital filters and generating TMS320 assembly code by converting the ASCII file containing the filter coefficients into fully commented assembly language code for all generations of the TMS320 family.
Price: $1350 U.S. for the ATD-C25; $750 for the ATD-FFT; $1500 for the ATD-FID.
Contact: ChiSoft Computer Institute, 9/7th Alley, Dong Shen Yuan, Wudaokou, Haidian District, Beijing, People's Republic of China, 886-2564527.
Inquiry 927.
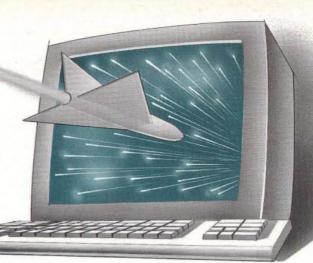
A Japanese-Language Oracle for the Mac

The Japanese-language version of the Oracle DBMS for the Macintosh supports HyperCard Structured Query Language inquiries and lets you generate I/O forms with HyperCard. Two versions are available: one where the database engine is local, and one for a network.

The network version, which is currently under development, will allow you to place the database engine in mainframes or Unix machines linked through datacom lines, letting the host handle the searches and updates and simply displaying the results on the Mac.
Price: 98,000 yen.
Contact: Ergonomics Software Products International Corp. (Ergosoft), Kiyou Building, Third Floor, 3-9-1 Akasaka, Minato-ku, Tokyo 107, Japan, 81-03-589-4951.
Inquiry 952.

Let Your PC Communicate with Transputer Systems
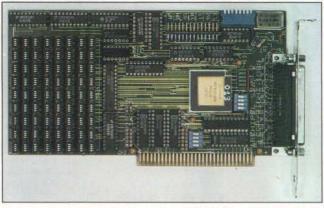
Archipel offers three boards that let you connect machines such as IBM's XT and AT, Apollo's DN 3x00, 4x00, and 100x00, and Sun's 386i workstation to transputer systems.

The Volvox-1/A incorporates a T414, T425, T800, or T805 transputer; 1, 2, 4, or 8 MB of DRAM; a C012 Inmos-compatible interface for communication between the PC bus and the transputer; DMA; and interrupts support. The transputer links are RS-422-level and are available on internal or external connectors. The full configuration possibility of the board is provided through jumpers, and the board's communication rate is more than 400K bytes per second.

The Volvox-1/AF adds a high-communication-rate interface between the PC bus and the transputer link that gives you a communication rate of more than 1 MBps. The price of this board is not yet available.

With the Volvox-1/V, you can connect VMEbus-based machines running under Unix or OS-9 with transputer-based systems. The board incorporates the same transputers, DRAM, C012 interface, and interrupts support as the Volvox-1/A, but the transputer's links are buffered through RS-422 differential drivers and are available only on external connectors.

You can control up to four autonomous subsystems by the transputer or the VMEbus. In addition, the Volvox-1/V can be a master or slave of the VMEbus. Two communication modes are available: One uses the C012 link adapter, and the other uses an original implementation of a

You can connect machines such as IBM's XT and AT, Apollo's DN 3x00, 4x00, and 100x00, and Sun's 386i workstation to transputer systems with the Volvox-1/A.

shared memory between the transputer and the VMEbus. The second mode offers a data or code transmission rate of 5.6 MBps at the application level.

You can program the Volvox systems in Occam, Parallel C, Parallel FORTRAN, and Parallel Pascal. The Volvox-Par tool offers easy integration of existing software in standard hardware and software environments.
Price: 18,140 French francs for the Volvox-1/A; 32,910 FF for the Volvox-1/V.
Contact: Archipel, P.A.E. du Levray, 4 route de Nanfray, Cran-Gevrier, 74000 Annecy, France, 33-50-69-20-58.
Inquiry 939.

A Fax Option for Dataflex Design's V.22bis Modems

Dataflex Design has added a fax/modem chip to its range of V.22bis modems that provides send-only fax capabilities. You can also retrofit the chip to existing Dataflex modems without any loss of data communications functionality.

BitFax, the fax transmission software, sends ASCII,

TIFF, PCX, and IMG files directly from disk to a Group 3 fax machine. The software's auto-dial directory employs the standard dBASE III file format, so you can use it to store other useful information on addressees; the directory/database is limited in size only by your machine's storage capacity. You can also combine files for transmission. Other features include a disk-based report and status log, broadcast functions, automatic redialing, and an automatic header line for your name and fax number.

Dataflex Design is bundling TransSend DataLink, new multilingual communications software based on Messiter Software's TransSend Plus package, with its entire range of Hayes-compatible modems. The software is available in 10 European languages and is suitable for virtually any asynchronous communications task. It employs user-definable pop-up menus and automatically generates script files by "learning" complete sequences of keystrokes. Script files let you log onto services with minimal effort and even permit automatic unattended operation.

The software features MNP level 5 error detection and an efficient encryption facility to protect sensitive data. In addition, it fully

supports most graphics standards (even in videotex mode) and offers complete emulation for the more popular terminals.

TransSend DataLink also features a simulated LED display that appears on the status line at the bottom of the screen and indicates TXD, RXD, DCE, DCD, and so on. You also get multitasking under Microsoft Windows/386 and DESQview; a DOS window; support for COM1, COM2, COM3, and COM4; on-line text and command-line editing; and XMODEM, ZMODEM, Sliding Windows Kermit, Kermit Server, CET software download, and proprietary file exchange protocols.
Price: £349 for the Biscom modems with fax option; £10 for TransSend DataLink for users of Dataflex Design modems.
Contact: Dataflex Design Ltd., Merton Park House, 2 Jubilee Way, South Wimbledon, London SW19 3XD, UK, 44-1-543-6417.
Inquiry 906.
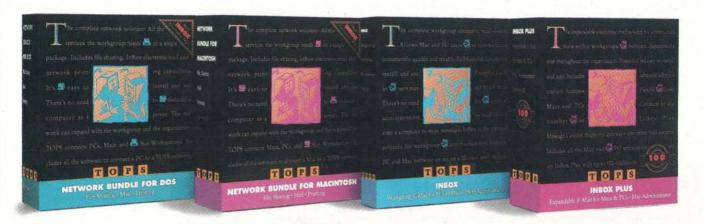
Programming with Helios

Helios-Sun provides a software development system with the Helios parallel-programming environment on Sun workstations with attached multiprocessor transputer systems. The system is a complete implementation of the Helios distributed operating system, fully integrated into the Sun networking and windowing environment.

You can run Helios over Ethernet networks from any workstation in a network of Suns, with Helios sessions appearing in multiple SunView

continued

Four reasons why 1990 will be the year of the workgroup.

It's already happening. More and more companies are embracing the idea of team or "workgroup" computing. Where people working on the same projects have the ability to exchange files, send electronic mail, share printers, and work together more effectively than even before.

Reasons enough for workgroup computing. But now TOPS presents four more...a new line of powerful networking products that will change the way you communicate from 1990 onward.

The ideal workgroup network— just add people.

With TOPS new Network Bundle for DOS and TOPS Network Bundle for Macintosh you'll have the power to share files, printers, and other resources with one person, your group, your department, or your entire organisation.

And best of all, both packages come bundled with the all new TOPS InBox, the powerful, yet easy-to-use electronic mail system. All, in one box.

Complete mail in a box.

On its own, electronic mail can be the key to a successful workgroup. And InBox provides fast, reliable, and easy-to-use electronic mail for up to 20 users. While InBox Plus delivers expandable communications for 100 users per Message Center.

In fact, used with other InBox Plus Message Centers, or UNIX Mail, MCI Mail, VMS, and other systems, InBox Plus can send messages anywhere in the world.

TOPS Network Bundle for DOS, Network Bundle for Mac, InBox, and InBox Plus. Four great products. Four great reasons to make 1990 the year of the workgroup in your company. And the best reasons of all to call TOPS today.

T O P S ®
When you need connections.

In the U.K., TOPS products are distributed by FrontLine (0256 463344), Go Mark (01-731 7930), and P & P (0706 217744).

windows alongside other SunOS applications. In addition, multiple users can share the transputer resources attached to the Sun network.

The basic software system consists of Helios, the Helios C compiler, and utilities. Options include FORTRAN 77, International Standards Organization Pascal, Modula-2, and a multiprocessor source-level debugger. Helios-Sun is available for Sun-3, Sun-4, and Sun386i workstations and supports most transputer boards for the Sun, including those from Inmos, Parsytec, and Transtech.

Distributed Software also offers the Helios PC Graphics Library, which you can use to develop parallel software on transputers that can use IBM PC graphics boards like CGA, EGA, and VGA. Helios pro-

grams call a unique graphics library, which is compatible with the Microsoft C graphics library provided on IBM PCs. You can then display the program output on your PC graphics system.

The Library is compatible with Microsoft C 5.0 graphics calls, so you can easily port applications you have written using Microsoft C to the transputer without having to change the graphical interface.

The Helios PC Graphics Library comes with source code and documentation. **Price:** £1750 and up for Helios-Sun; £300 for the Helios PC Graphics Library. **Contact:** Distributed Software Ltd., 670 Aztec West, Almondsbury, Bristol BS12 4SD, UK, 44-454-612777. **Inquiry 907.**

Fela's Laserplotter for Printed Circuit Board Makers

The HoloFlash Laserplotter gives small or medium-size manufacturers of printed circuit boards (PCBs) the highest standard of film production and quality in-house, according to Fela. When you use it with an IBM PC, XT, or AT, HoloFlash forms a stand-alone or network component for CAM systems.

The flatbed plotting surface measures 510 by 660 mm, and you can set the resolution at 12 microns (2000 dpi) or 6 microns (4000 dpi). Regardless of the complexity of the circuits, the plotting time for the entire surface is 4 minutes

at 2000 dpi.

The FlashCAM software is designed for high-precision PCB photoplotting. Its user-friendly menu system simplifies the input of CAD data and the processing of Gerber, Quest, Hewlett-Packard Graphics Language, and numerical control drill data.

You get multilayer viewing (up to 35 layers) with independent data formats and aperture displays, WYSIWYG custom apertures, design rule check, true rotation of asymmetrical pads, step and repeat, and pattern generation for thieving and venting. **Price:** 210,000 Swiss francs. **Contact:** Fela Tec AG, Rue Marziano 21, CH-1211 Geneva, Switzerland, 41-22-43-85-53. **Inquiry 945.**

A Little Something Extra for Heavy Users

Time and Space-

the ARC Accent SX gives you more of both.

Maybe you need more time at home away from the office but need access to a PC. Or, maybe the office is a small area in your home with no space for a conventional PC.

The Accent SX is so easily portable. Just close it up and take it with you. You can set it on your lap and work comfortably in any room in your house. And you can powerfully automate the home office in as little space as is required to open your brief case.

Accent on Power

The Accent SX is small but only in stature. Performance-wise it makes a lot of conventional PCs seem like slow motion Behemoths. It's motivated by an 80386SX-16 CPU with a turbo page mode speed up to 19.4MHz. It comes with an internal 40MB hard disk drive with an access time of 28ms. It also has 1MB of RAM, expandable to 4MB. You could search for a long time and not find an application that will slow the Accent SX to less than light speed.

A Herd of Features

Few full size monitors offer a view as good as the Accent's Gas Plasma VGA screen which tilts 130 degrees to your most comfortable viewing angle. Along with standard features such as a 1.44MB internal Floppy Disk Drive, 80387-16 Math Co-processor socket, 84+

introducing a little something extra, The ACCENT SX

key Multi-function Enhanced Keyboard and a Real Time Clock/Calendar, interface capabilities are limited only by your imagination and:

2 RS-232C Serial Ports
1 Parallel Printer Port
1 16-bit AT Expansion Slot
Universal Monitor Port

And it all fits in a 376 x 345 x 96mm case and weighs only 7 KGS.

The Accent SX from American Research Corporation - a little something extra that fits the biggest users and the smallest places.

ELEPHANTS IN BUSINESS

accent sx

ARC
AMERICAN RESEARCH CORPORATION

Upcoming
International Events

INNC-90

The International Neural
Network Conference will
be held on July 9–13 at the
Palais des Congrès in Paris,
France. The event will in-
clude one day of tutorials de-
signed to introduce the tech-
niques, developments, and
prospects of neural networks
to nonspecialists; four days of
conferences chaired by 200
of the world's top scientists and
industrialists in the field; an
overview of commercial activi-
ties in the U.S., Europe, and
Asia; the presentation of prod-
ucts and prototypes; and a fo-
rum with workshop sessions.
Contact: Nina Thellier
N.T.C., INNC-90-Paris, 19,
rue de la Tour, 75116 Paris,
France, 33-1-45-25-65-65.
Inquiry 957.

OIS
International '90

Europe's leading informa-
tion and image manage-
ment conference and exhibi-
tion, Optical Information
Systems International '90,
will be held July 17–19 at the
Wembley Conference Centre
in London. The range of topics
to be covered includes docu-
ment image processing, optical
data storage, full text retriev-
al, electronic document man-
agement, CD-ROM publish-
ing, drawing management, and
micrographics systems.

The event is designed to
provide attendees with guid-
ance in areas such as the for-
mulation of an integrated
information and records
management strategy for an or-
ganization or division, an
analysis of the key business
benefits that document
image-processing and docu-
ment management systems
can deliver, first-hand ac-
counts by users, a differenti-
ation of the requirements
placed on the technology by
the major application areas,
and a comparison of the dif-
ferent approaches that suppli-
ers are taking to solve docu-
ment management, processing,
and distribution problems. In
addition, there will be an expo-
sition of the significant new
developments over the last
year, an analysis of the key
technical issues still facing the
industry, and a review of the
current position on standards.
Contact: Meckler Ltd.,
Grosvenor Gardens House,
Grosvenor Gardens, London
SW1W 0BS, UK, 44-1-931-
9985.
Inquiry 958.

Iberamia-90

The second Ibero-Ameri-
can Conference on Artifi-
cial Intelligence is scheduled
for July 9–13 in Mexico.

The objective of the con-
ference is to promote the inter-
change of ideas among Iberi-
an and Latin American
researchers working in AI,
while encouraging active par-
ticipation from all sectors of
the international AI
community.
Contact: Iberamia-90, Attn:
Srita. Ma. Antonieta Alvarez
Pérez, Apartado Postal
70302, C.P. 04510, Mexico,
D.F., 52-554-4723.
Inquiry 959.

Upcoming International Events

Hong Kong Computer Expo '90

The Sixth Hong Kong Computer Expo, incorporating CAD/CAM '90, will be held on May 7–10 at the Hong Kong Convention & Exhibition Centre.

The show will feature computer hardware from mainframes to personal computers, peripherals, data communications, add-ons, accessories, and supplies, while the CAD/CAM section will showcase the latest technology and its diversified applications for industry. The Second International Software Show for Asia, called Software Village '90, will run concurrently in the Centre. **Contact:** Business & Industrial Trade Fairs Ltd., 28th Floor, Harbour Centre, 25 Harbour Rd., Wanchai, Hong Kong, 852-5-756333. **Inquiry 953.**

Real-Time Systems Workshop

Euromicro has announced that real-time systems will be the subject for a workshop on June 6–8 in Horsholm, Denmark. The workshop will cover topics such as real-time architectures, real-time applications in industrial control, testing and fault tolerance, real-time communication networks, CAD for real-time systems, real-time programming and languages, AI application in real-time environments, and scheduling. **Contact:** Mr. J. Molgaard, Organizing Chairman, Elec-

tronikcentralen, Velinghedsvej 4, 2970 Horsholm, Denmark, 45-42-867722. **Inquiry 955.**

Euromicro '90

The Sixteenth Symposium on Microprocessing and Microprogramming, or Euromicro '90, will be held on August 27–30 in Amsterdam, The Netherlands. The theme of the event is Hardware and Software in System Engineering with an emphasis on the hardware/software interface and system reliability.

Topics will include hardware and architectures (e.g., RISC architectures and design issues, VLSI algorithms and architectures, application-specific integrated circuits, and special architectures), software (e.g., software engineering, operating systems and development environments, and software reliability for real-time systems), systems (e.g., fault tolerance and diagnosis; computer-integrated manufacturing, components, and techniques; process control and real time; and computer networks and internetworking), and applications (e.g., advanced applications, innovative parallel applications, image and vision systems engineering, and application systems reliability). **Contact:** Euromicro, Attn: Mrs. Chiquita Snippe-Marlisa, P.O. Box 545, 7500 AM Enschede, The Netherlands, 31-53-338799. **Inquiry 960.**

ROCK AND REDEEMER

You deal with enough PC parts and systems vendors to know that for most, it's hard to find many redeeming features in their supply programs. At AGC Technology Corporation, we know dealers and distributors are people too. Which is why we have created a solid personalized and flexible distributorship program that caters to your needs and concerns.

When it comes to customer service, AGC reliability means we come to you. We have well-situated offices in the U.S.A., Spain and West Germany to provide on-hand technical and sales support. Our representatives regularly visit customers in their own places of business for convenient, personal, face-to-face service.
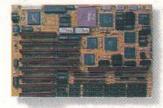
To make sure our communication is as clear as possible, we have staff fluent in German, Spanish, French and Italian, as well as in English. So we can listen to your concerns, understand them and respond.

In any language, though, excellent quality is easily understood. AGC produces a full scope of skillfully designed VLSI, NEAT, 386SX, 386/AT and 486/AT main boards for sale alone or as engines for our fully integrated AGC systems.

Shipment schedules are made for a reason, and we stick to them. No excuses. Our delivery system is fast, efficient and always on time.

If you're ready to be treated on a higher level than you presently are. If all you really want are solidly engineered products, backed by responsible service at a competitive price. Then turn to AGC. We just may be your rock and redeemer.

Upcoming International Events

TOOLS '90

The Second International Conference and Exhibition on the Technology of Object-Oriented Languages and Systems (TOOLS) '90 is scheduled for June 26–29 in the CNIT Center in Paris (La Défense), France.

The event will include tutorials; invited presentations by international object-oriented experts; submitted papers on important practical aspects of object-oriented techniques; workshops; and demonstrations of object-oriented tools, languages, environments, databases, and their applications.

Contact: TOOLS '90, SOL, 14, rue Jean Rey, 75015 Paris, France, 33-1-40-56-03-58. **Inquiry 961.**

ISS '90

The Stockholm International Fairs and Congress Centre in Alvsjö, a suburb of Stockholm, Sweden, will be the location of the International Switching Symposium (ISS), which will be held on May 27–June 1.

A specialist conference in the telecommunications field, ISS has been held every three years since 1957. It is an international technical scientific conference attended by producers, operators, and administrators in the telecommunications field.

The topics being presented will include broadband network evolution, signaling and protocols, optical switching, voice services, ISDN evolution, object-oriented software, metropolitan area networks, software development and portability, ATM switch architecture, evolution of intelligent networks, testing and quality, network performance, rural networks, Telecom management networks, and mobile and satellite communication.
Contact: ISS '90, c/o Congrex, P.O. Box 5619, S-114 86, Stockholm, Sweden, 46-8-6126900.
Inquiry 954.

MacWorld Expo Asia '90

MacWorld Expo Asia '90 is scheduled for June 14–17 at the Raffles City Convention Centre in Singapore.

Exhibitors at the show will include Apple, Aldus, Ashton-Tate, ACI, SuperMac, Microsoft, Tektronix, Hewlett-Packard, Digitalk, Micro Planning International, and CMS.
Contact: MacWorld Expo Asia Pte. Ltd., 39 Sixth Ave. Level 2, Guthrie Building, Singapore 1027, 65-337-2604.
Inquiry 956.

continued

Develop Neural Networks

You can develop neural networks with Neurun, an IBM PC–based neural network development system using the McCleland and Rumelhart back-propagation algorithm, the Quickprop algorithm, and the Solis and Wets algorithm.

With this system, you can generate networks of arbitrary size and complexity. Industry-standard interfaces allow you to easily import training data and export results. You can also use Neurun as a development tool and an evaluation aid.

The network developer designs a network using NCL, Neural Computer Systems' proprietary language, and either includes the training pattern set in the source file or reads it from a separate file at run time. You can include the input data in the source file or load it from a Lotus 1-2-3–compatible spreadsheet. The Neurun system then runs the network defined, draws the network on a graphics screen, and provides a strip graph of the progress of the network. You can also output network solutions in Lotus 1-2-3 format and generate a log of the network state as the training progresses.

Neurun requires an IBM PC with 512K bytes of RAM, DOS 3.0, and a Hercules, CGA, EGA, or VGA graphics card. The company recommends a hard disk drive and an arithmetic coprocessor.

The company also offers Neurun Light, a junior version of Neurun that implements the back-propagation algorithm, which has emerged as the most popular choice for research and practical applications. Neurun Light uses a subset of the NCL language to define the network, the learning parameters, and the data to be learned. It presents you with a graphics display showing a diagram of the network and a strip chart showing its progress. In addition, Neurun Light is upgradable to Neurun at any time—the hardware requirements are the same.

Price: £450; £99 for Neurun Light.
Contact: Neural Computer Systems Ltd., 79 Olney Rd., Emberton, Olney, Buckinghamshire MK46 5BU, UK, 44-234-713298.
Inquiry 938.
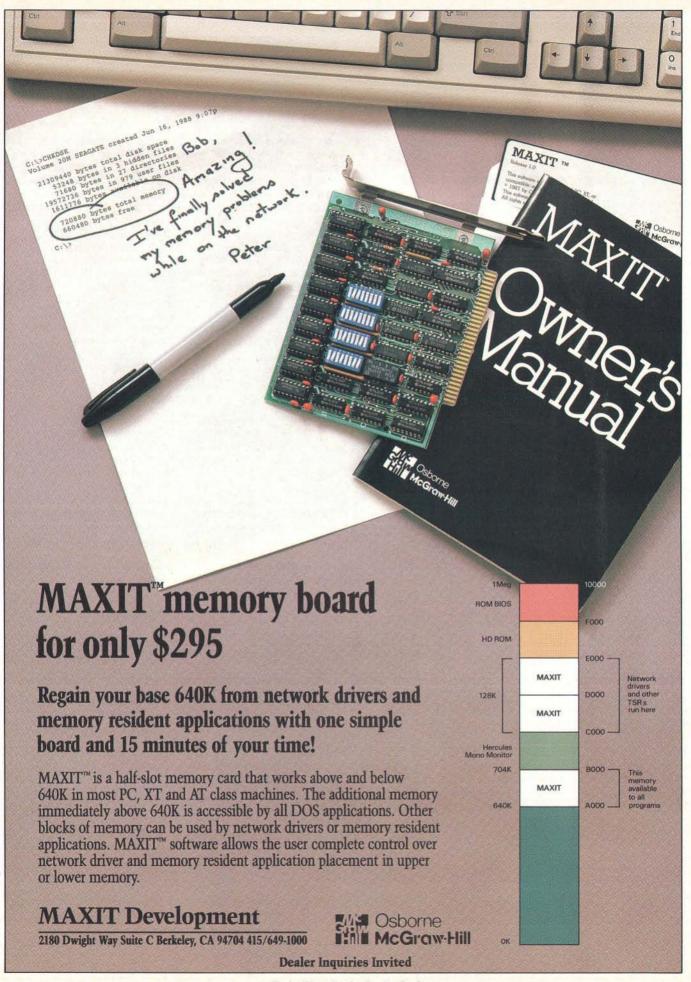
Software Control for Genicom's 3820 Printer

Genicom's 3820 high-speed, heavy-duty, dot-matrix printer lets you change paper paths and printer emulations through software. Four paper paths are available, including continuous-form paper feeders at the bottom, rear, and front of the machine, and a top manual cut-sheet paper feeder. You can print alternate forms on the three continuous-path feeders without unloading or reloading paper. In addition, the printer features automatic paper parking and zero-inch forms tear-off.

Emulations for the model 3820 include ANSI X3.64, which lets you take advantage of oversize and bar code features, plus IBM Proprinter and Epson. The printer comes with an RS-232C serial interface, and you can use it in a LAN.

The 3820 prints 600 cps at a resolution of 12 cpi and 500 cps at 10 cpi in draft mode. Letter-quality mode gives you 150 cps at 12 cpi and 125 cps at 10 cpi.

Genicom is also releasing a pedestal for the 3820 that

Circle 451 on Reader Service Card

accommodates two boxes of stationery for input to the printer via the bottom, rear, or front paper paths and that features a shelf to capture and store printed forms.
Price: £1995; £2195 for the bar code version.
Contact: Genicom Ltd., Armstrong Mall, Southwood, Farnborough, Hampshire GU14 0NR, UK, 44-252-522500.
Inquiry 934.

A Universal In-Circuit EPROM Emulator Card

PCEMU is Alfatron's universal in-circuit EPROM emulator card for the IBM PC, XT, and AT that supports 8-, 16-, and 32-bit systems and 27(C)64, 27(C)128, 27(C)256, 27(C)512, 27C1001 (1-Mb), and 27C2002 (2-Mb) EPROMs.

The card is suitable for embedded development systems. It can virtually eliminate the step of burning EPROMs from the design cycle and requires only a small amount of time to download the code, according to the company. The PCEMU is completely transparent to the target system and allows true execution of system software in the target system hardware environment.

You can use the card to emulate popular 28-, 32-, and 40-pin 200-ns EPROMs and PROMs from 8K bytes up to 256K bytes (2 Mb) in size. It can emulate a single 8-bit-wide EPROM, a pair of 8-bit-wide EPROMs in 16-bit-wide configuration, or a single 40-pin, 16-bit-wide EPROM. In addition, all these features are programmable. If you want to use the PCEMU in 32-bit systems, you will need two cards.

A programmable parity open collector driver lets you connect to the target system's reset line so that you can hold the target CPU in reset state while the host PC loads new data into the emulation memory. As soon as the loading is complete, the reset line releases and the target CPU starts executing the program.

The PCEMU software driver supports Extended Intellec-86 hexadecimal download format, Motorola-S download format, and Binary Image download format. A simple command-line interface lets you configure the PCEMU according to the target system's requirements. In addition, you can include the PCEMU command line in your project's batch or make file to automate the download and run cycle.

PCEMU requires 256K bytes of RAM and DOS 3.0.
Price: $A 800.
Contact: Alfatron Pty. Ltd., 5/14 Jersey Rd., Bayswater, Melbourne, Victoria 3153, Australia, 61-3-7205411.
Inquiry 940.

Generate C Programs

System C's latest program generator, called Sycero C, is aimed at two types of users: experienced C programmers who would like to write programs more quickly, and developers who are new to C but need to become productive in a short amount of time.

Sycero C is based on a data dictionary that lets you define an unlimited number of indexed files per system. It supports standard field types, including numerics, dates, strings, and variable-length memo fields.

You can paint screens and include fields on each screen from up to 99 files. Sycero C supports lines, boxes, full

continued

BIX IS GLOBAL

BIX is BYTE's Information Exchange, a worldwide computer conferencing system devoted to the exchange of microcomputer information.

When you use BIX, you leap borders of geography and time to exchange information, opinions and ideas with a "living database" of the world's most knowledgeable microcomputer users.

BIX covers the world.
Our *Microbytes* news service, backed by BYTE and McGraw-Hill, provides worldwide daily news coverage of major events in the microcomputer industry plus listings of thousands of new products from vendors around the globe. Scores of companies now use BIX to provide technical support for their hardware and software products worldwide.

BIX saves you time.
BIX resources help you get the micro-related information you need right away, regardless of your location. BIX electronic mail lets you send private messages to other BIX users worldwide. For more information on how to join and use BIX from your country, read the BIX message in this issue (see Advertiser Index for page number), or contact us today.

BIX
BYTE INFORMATION EXCHANGE

One Phoenix Mill Lane
Peterborough, NH 03458 USA
(603) 924-7681 (Our overseas
helpline number)
8:30 AM–11:00 PM
Eastern Time
(–5 GMT) Weekdays

color, and windowing. The screen handler also includes a word processor that lets you edit text in a window on the screen.

Sycero C also includes a report generator that automates up to 10 levels of subtotal, sort sequence selection, report ranges, wild-card matching, and exception reporting. You can paint report layouts, and each report can contain information from up to 99 files.

Based on dBASE, the Sycero Programming Language lets you use dBASE syntax and generate C. SPL supports more than 170 dBASE and Clipper commands and functions, and more than 200 of its own. It also provides a fully integrated library of file, screen, printer, and variable handling functions.
Price: £595.
Contact: System C, 60–61 High St., Maidstone, Kent ME14 1SR, UK, 44-622-691616.
Inquiry 933.

The Latest Version of X.desktop

The latest version of X.desktop adopts the look and feel of Motif through the use of the standard Motif toolkit and features professionally designed three-dimensional icons.

X.desktop 2.0 provides an icon-based graphical representation of the files and programs on a computer or network. Instead of resorting to the Unix command line, you can run programs and carry out file management activities such as copying, deleting, creating new files and directories, printing, and archiving

by simply manipulating these icons.

The new version also features a supply cabinet facility for individual users, giving them rapid access to widely used programs, files, and directories on a network. The additional configurability of the menus allows a system administrator to add, delete, and change options from any X.desktop menu. For example, if a system administrator doesn't want users to have access to the Unix shell, he or she can tailor the menu to remove that option from the menu altogether. Version 2.0's extended user-friendly rule language lets you use full-length words instead of abbreviated codes when you are editing X.desktop rule files.

X.desktop 2.0 runs on the Sun-3, Sun-4, Apollo Domain, Hewlett-Packard 9000, Sony NeWS, Mac II, DECstation, and Intel 386–based personal computers. It conforms to the X/Open Native Language System standard, which lets you easily translate the software into different languages and dialects. In addition, the software supports double-byte character sets for easy translation into Asian languages.
Price: £395.
Contact: IXI Ltd., 62–74 Burleigh St., Cambridge CB1 1OJ, UK, 44-223-462131.
Inquiry 924.

Fix

Scientific Text Processing informs us that the price of its ChiWriter, a scientific, multifont word processor (see January What's New International, page 80IS-34), should be £109.95.
Contact: Scientific Text Processing Ltd., 53 Westridge Rd., Portswood, Southampton SO2 1HN, UK, 44-703-557632.

The Cream.   The Crop.

There are plenty of places to get information in this industry. Too many. But if you want the best quality information, there's only one that rises to the top: BYTE*WEEK*.

BYTE*WEEK* is a weekly newsletter from the same professionals who produce BYTE Magazine. Each week, the most important news and information from the previous week is presented in a readable and concise manner. BYTE*WEEK* offers you what no other publication can: timely news on the rapidly-evolving computer industry as it happens with the interpretation and evaluation that only BYTE's experienced editorial staff can provide.

Subscribe now and take advantage of a special subscription rate of $395 ($495 outside the U.S. and Canada). Your subscription to BYTE*WEEK* also includes a free subscription to BIX, BYTE's exclusive on-line conferencing system. *Don't miss this opportunity!*

For fastest service, call toll-free **1-800-258-5485** (in N.H., call **603-924-9281**) and charge to a major credit card or we'll bill you.

BYTE*WEEK*
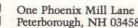
One Phoenix Mill Lane, Peterborough, NH 03458.

BYTE*WEEK* offers a money-back guarantee if you are not completely satisfied.

BACKUPS, FAX, AND MAC DISASTERS

The mail brings some interesting things, including globes, teddy bears, and a pair of CD-ROM drives

I t all started when I wanted to check out a new CD-ROM for the Macintosh. My Big Mac IIx is to the right of my desk. I have a small keyboard table for the Mac, and when really interesting software for it comes in, I generally pile it there.

You turn on the Mac by pushing a button on the keyboard. I did that and turned back to Big Cheetah. After a while, I noticed my Mac wasn't playing his start-up music, and when I looked around, he was displaying AASK, which is an initialization utility manager. AASK comes to the foreground if you hold down the space bar while the system is starting up. Sure enough, some of the software on the keyboard table was holding down the space bar and a couple of other keys.

I cleared the clutter off the table. Then I must have taken leave of my senses. Instead of reaching around to the far side of the Mac and pushing the "programmer's button"—hardware reset for the rest of us—I simply turned the Mac off.

I turned the Mac back on. It wouldn't boot. It did strange things, flashing stuff up and taking it down too fast for me to see, and eventually it just sat there with a picture of a floppy disk and a question mark. I used hardware reset. Same thing. Again, and again. I turned it off at the switch, waited a while, and turned it back on. Same thing. That Mac wasn't going to boot nohow.

I found a System floppy disk and booted with that. The Mac came up all right, but it erroneously reported that my write-protected System floppy disk was

damaged and offered to initialize it. I quickly canceled that offer.

The Mac showed only two icons: the utility floppy disk and the big Priam 330-megabyte MacDisk. No sign of the 80-MB internal hard disk drive.

Don't panic, thought I. Try the program known as First Aid. That didn't work. First Aid thought there was an unnamed disk in addition to the floppy disk and the MacDisk, but it wouldn't do anything about it. Same thing. I put a System folder in the MacDisk and rebooted. The Mac told me my hard disk was damaged and offered to initialize it. I clicked cancel with a shudder.

By now I wasn't in a panic, but I was pretty upset. Who, I thought, who do I know who knows Macs, and I know well enough to call at 10:00 on a Sunday night? I certainly couldn't call BYTE's Tom Thompson; if it's 10:00 p.m. in Hollywood, it's 1:00 a.m. in Peterborough. Then I thought of Alex's friend Stefan Somogyi, who knows Macs better than anybody except maybe Tom. Stefan recently returned from a trip to Berlin, where he grew up, so it would be interesting to talk to him anyway.

Fortunately he was home, and I went through the symptoms. Stefan thought for a moment. "I never saw one do that before."

Now, I thought. Now it's definitely time to panic.

"Boot up with the floppy disk again," he told me. I did. No sign of the icon for the internal hard disk drive. "Have you got the First Aid program on that floppy disk?" he asked. I did. "Good. Pull the plug on the MacDisk, and boot up with the floppy disk."

I did that and launched First Aid. Then I clicked on the Drive button in the First Aid display; and lo!, it had found the 80-MB hard disk drive. I told Stefan.

"Great. Now before you tell First Aid to repair that disk, do a clover-S." I dutifully pressed the clover—I call it the splat—key and the S key. A new small

window appeared on-screen. "Now tell First Aid to work on it."

It took a few minutes, but by gollies that did the trick. My 80-MB internal disk drive was back in action. Not only could the machine find it, but it could boot from it as well. The next thing to do was to plug in the MacDisk and reset while holding down the clover and option keys. That rebuilds the Desktop. With that done, all is well.

The moral of the story is that Mac software isn't as robust as DOS software. Apparently you should never turn off your Mac as a means of shutting it down: which leaves you in an interesting situation if you have a power failure and no uninterruptible power supply. I already have mine on a surge suppressor; now I'm going to connect it to the Clary UPS as well.

PC Librarian

If you have a properly installed and managed network, the Palindrome system that makes automatic archives to tape cartridges (see my April column) works wonders; but most of us don't have anything that fancy and have to be satisfied with something less: floppy disks for the most part, although in my judgment if your work is worth much, you'll invest in a WORM (write once, read many times) or a Bernoulli Box.

Either works pretty well. You use the DOS XCOPY command with a specified date, which copies everything accessed since then, and all's well. The WORM drive keeps everything sent to it, so that you can page back through and find previous versions of your work, which is a good thing: one corollary of Murphy's Law is that you will never need a previous version of your work until the day after you have erased the last copy.

You can do something of the sort with the Bernoulli Box by saving off various versions under different names or archiving Bernoulli disk cartridges. The

Bernoulli disk cartridges don't hold as much as the WORM glass disks do (over 380 MB per side in the Maximum Storage WORM I'm using), so you have to do your own management and record keeping, which can get pretty tedious.

In addition to the backup problem, most of us need actual archiving: storing programs and files where you can get at them again, but out of the way so they don't clutter up your hard disk. That's where PC Librarian comes in.

PC Librarian will let you mark entire directories (with or without the subdirectories) and copy all those files into one large archive file, compressing them if you like. Another option is to encrypt them. PC Librarian builds a catalog of files you have archived; the catalog not only tells you where you put what, and when, but also lists the date, time, attributes, and file size for each file. You can print that catalog or view it on-screen.

Finally, it will optionally remove all archived files from your hard disk after it has copied and verified them.

The archive files so created can be in another directory on your hard disk; on a WORM drive; on a Bernoulli removable

disk cartridge; on floppy disks; or on some distant medium accessed through a network. It does all this neatly and with little effort on your part. That's the good news.

The bad news is that it's slow when writing to floppy disks. Archiving 12 MB to 3½-inch 1.4-MB floppy disks (it needed 8 MB to store everything in compressed form) took something over an hour. However, archiving to a hard disk or WORM drive is much faster—about as fast as XCOPY provided that you're not compressing and encrypting.

There are some idiosyncrasies. For one, if you want to store on a WORM, you have to tell the system that's a "virtual drive"; unfortunately, the manual doesn't tell you that. It does insist that you use the Install program, which expands a number of the files. Clearly, they have archived much of the program itself in compressed form. I found that the best way to install the program was to copy the files, including the Install program, from the distribution floppy disk to a \PCLIB subdirectory on my hard disk and then run Install from there.

Once you've got it installed and have

worked with it for a few minutes—I advise you to practice archiving files you don't care about until you get the hang of it—it becomes reasonably intuitive. The manual is complete and readable, but you don't have to read it except to search for features. There's a lot of good on-line instruction to handle most questions.

One good thing about PC Librarian is that the archive files and catalogs it makes are ordinary DOS files. That is, although you can't *read* the archive files (at least not in compressed form), you can copy the archive file along with its catalog and index files. In my case, when one archive got a bit large, I stored it and its associated files in a subdirectory named \PCLIB1 on a WORM drive and started a new one. I now have about four of them. To retrieve those files, I need only copy the archive and its associated index and data files back to the active area on my hard disk; PC Librarian does the rest.

The only real problem I've had with PC Librarian is having to run the Rebuild Catalog program a few times when I didn't think I should have to do that. In

every case, though, it worked fine, and no files were lost.

PC Librarian will tell you where your files are (by volume name if they're on floppy disks), when you did archiving sessions, and what you did during them. In fact, it generally does all the things you wish your librarian program would do. I like it enough to give it a User's Choice Award.

On that score, last month I ran out of space before I had finished with my

awards; there really were a lot of good developments last year.

Northgate
Northgate Computer Systems burst on the scene in full force last year, with a full line of well-made and smooth-running computers, including fast 386s. They also make my favorite keyboard; and they distribute, at their cost, POURVOUS.COM, a tiny (20-byte) TSR program that makes the comma and period

keys print commas and periods shifted or unshifted. (You do Alt-Shift to get the < > characters.)

The really outstanding feature about Northgate has been their technical support, which has become the standard that other mail-order computer companies must match. I have no problem giving Northgate two User's Choice Awards: for their keyboards, and for their technical-support group.

Real Science
It's not strictly a microcomputer topic, but in my judgment one of the most successful government agencies is the National Science Foundation, which supports fundamental science research and education. The National Supercomputing Centers supported by NSF have been spectacularly successful.

Designed to give American students a chance to familiarize themselves with supercomputers while in college, these centers quietly turn out new software and develop new ways to interface humans to computers—as well as to interface microcomputers to supercomputers.
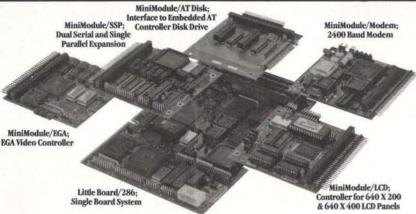
The Chaos Manor Government Project of the Year Award goes to the NSF for its support of the National Supercomputing Centers, and I sure hope Congress not only pays attention, but doubles the budget. National support of the sciences receives a trivial amount for what we get for it. And now that Europe is reuniting, it's vital to give our students every advantage possible. If we're going to borrow money from our children, doesn't it make sense to invest some of it in their future?

OS/2
I don't recommend that users get OS/2 just yet. It has a lot of good features, but it also has the near-fatal flaw of requiring you to buy all new software. It's true that the OS/2 version of a program will likely be better than the DOS version (with some notable exceptions), but it's also true that for many of us, our DOS word processors and spreadsheets and CAD programs are good enough. Still, some applications actually need enormous blocks of memory and run smoother under OS/2 than under DOS with Phar Lap Software or Rational Systems extensions. Some people need OS/2.

If you're one of those, you have three choices: do without your DOS applications, run two machines, or use Polymod 2, available from The Software Lifeline. With Polymod 2 you'll still need two computers, but you can control one with

continued

the other; that is, you set up your DOS system as the master and then run OS/2 on a fast 386. You can then control the OS/2 machine with the DOS system, meaning that you can start processes in OS/2 and leave them running while you go about your normal DOS business. You can even run Polymod 2 in a DESQview window, which is a bit weird, but it works. I presume it would work with Windows as well, although I haven't tried that. What you must do is run a terminal-emulation program on your PC.

It's also possible to connect a number of terminals—up to eight—to the OS/2 system. I haven't done that, but it doesn't look as if it would be all that difficult. This makes your OS/2 machine a multi-user, multitasking workstation. If all terminals are actually PCs, you have at least one CPU per user, and I can think of circumstances in which you'd want to have multiple users on an OS/2 system; but I sure wouldn't load eight users on it at the same time.

Polymod 2 has password protection and what appear to be fairly good security procedures. All in all, it's a reasonable solution for a small business office consisting of three employees, each with a PC, and perhaps a couple of extra terminals used intermittently.

Now a warning: I am no OS/2 expert. I managed to get Polymod 2 set up so that I was controlling OS/2 on a 33-MHz 386 through Big Cheetah, my 25-MHz 386 running DOS 3.3. The Polymod 2 documents are good enough to accomplish that; and I had some help, in that I was sent a terminal board for my PC and a

cable set. All the cabling information (it needs at least seven wires of an RS-232C cable wired exactly to a given specification) is in the manual. A breakout box helps a lot.

Once it was set up, I was able to do some OS/2 things, including running the Logitech Debugger, with the system. Everything seemed to work all right. I then took the system down, because I have no real need for it and an associate badly needed to borrow the OS/2 machine. I can't claim, then, to have given this a thorough test, nor am I really the right person to do that kind of test. However, it's my opinion that this may well be a relatively painless way to go from DOS to OS/2, since you can gradually move more and more applications to OS/2 while not yet abandoning DOS.

One final warning: Polymod 2 employs a copy-protection scheme that is a bit strange. Sometimes on start-up it asks you a question. There are little colored blocks at the bottom of each page of the manual, and each block has three alphanumeric characters printed in it; you must read the code off the appropriately colored block from the proper page or the system won't work. It's therefore guaranteed to drive color-blind users out of their minds. I don't have any trouble with the scheme, but I'm not color-blind.

Things My Postman Brings Me
First it was globes. Little ones on key chains. Slightly larger ones embedded in Lucite. One the size of an orange on a stand. Another the size of a grapefruit. They came in unmarked packages, with

no return address, with banner print proclaiming THE WORLD IS COMING.

Eventually it did: the World Factbook 1989 Database on a CD-ROM. This is the unclassified database prepared annually by the CIA. It's available from Quanta Press, and it now works on both PCs and Macs. What they sent was the software for the Mac and the CD-ROM—indeed, the 3½-inch Mac disk was inside the CD-ROM case. This is the disk I was trying to look at when I had the Mac accident I mentioned earlier. Fortunately, the software for reading the CD-ROM on the PC is on the CD-ROM itself.

I've talked about this product before. It's a reasonable database, although not as complete as I would have thought. I presume the Mac version is identical.

Then came another large box. This one contained a teddy bear dressed in a camouflage uniform, complete with Green Beret, MacV patches, jump boots, and an extraction harness. A snappy-looking fellow, actually, with a marked resemblance to the colonel (now general) who invented the Green Berets.

There was also a military-style belt pouch containing a MacV SOG (Military Assistance Command, Vietnam, Special Operations Group) bandanna, a pencil, and a CD-ROM that claims it's a database about Vietnam.

I say claims because, while I can get a directory from the CD-ROM and install the software that's supposed to read it, whenever I try to access the thing, I get messages saying that the search has failed. That's on Big Cheetah. Trying it

Microsoft creates software giants in laboratory.

A class at Microsoft® University will go straight to your head.

Reason being, your course instructors work for us, Microsoft. The country's leading developer of software.

Better still, they do their teaching in a laboratory setting that gives you two big advantages: Hands-on experience. And software you've developed that's yours to keep.

It's the fastest way to learn the latest technology being utilized in today's popular programs. Giving you, and your corporation, a big jump in developing software.

Courses are offered in several power-ful systems platforms, including Microsoft OS/2, Microsoft OS/2 Presentation Manager and Windows.™ And innovative networking technologies like LAN Manager and Microsoft SQL Server.

To get more information and a free copy of the Microsoft University catalog, call (800) 541-1261, Dept. K88.

As a graduate, you'll soon be writing better applications, faster.

Making your career, and your company, grow by leaps and bounds.

MicrosoftUniversity™

on the Zenith Z-386/25 with a Tecmar VGA board gets me the message that the program doesn't support my graphics board.

Unfortunately, this isn't much of a surprise. For reasons I don't understand, no two publishers of CD-ROM software for PCompatibles use the same access programs, and about half of them don't work very well. I thought for a while I had done something wrong in the installation, but that wasn't so. I followed the instructions, and I do get into the database. It just won't read anything on either machine.

I really like the teddy bear, but I sure wish people would pay as much attention to making their software work as they do to calling attention to it.

Disk Utility Blues

Another recent arrival contained coffee mugs with CHAOS MANOR glazed onto them along with Golden Bow's Cupid logo. They're my favorite cups, duplicates of the ones Golden Bow sent last year. I'm glad to have got them, and I do use and recommend Golden Bow products, whether or not they send me coffee cups; but, alas, their disk utilities won't work with Roberta's Perstor hard disk drive controller (see my March column). Incidentally, that controller works quite well: fragmented files or not, it's a *lot* faster than her old controller was.

Actually, almost everyone's disk utilities are behind the times. As an example, I tried OPTune on the Premier 9000 (a 33-MHz 386). It won't work. The Premier has a disk cache on the motherboard, and if there's any way to turn it off, I haven't been able to find it. In addition, while OPTune (and Golden Bow's Vopt, for that matter) will repack the files on the disk, it won't do any other tune-up operation.

Neither will Gibson's SpinRite II: it too disqualifies itself when I try to run it on the Premier, even though the SpinRite package *says* it will work with DOS 4.01 and large disk drives. It sure won't with this one. It does partly run on the Z-386/25 (which runs a special Zenith version of MS-DOS 3.3 so that it can support a 150-MB drive), but even there it won't do its primary function. Neither will OPTune.

OPTune, SpinRite, and other such programs dynamically reformat your hard disk; that is, they pick the data off a sector, store it, reformat that sector, and put the data back down. SpinRite used to have the problem that it would try to guess if it found an ambiguous spot on your disk, which is all right, but it did it without even *telling* you, and that's not acceptable. The new version doesn't have that problem. If you have a standard machine, you can and should use one of these programs. I just wish they'd get versions I can use on really high-performance machinery.

Denon CD-ROM

For the last few years, I've used Amdek Laserdek CD-ROM drives with Big Cheetah. The drives are manufactured by Hitachi, and I've never had any trouble with them, barring that any CD-ROM drive eats memory. If you have a 386 and use a memory manager like Quarterdeck's QEMM-386, you can load the CD-ROM access software into that high memory area between 640K bytes and 1 MB before loading DESQview.

Then the CD-ROM drive is available in any DESQview window; as an alternative, you can load the .SYS portion into high memory and run the MSCDEX .EXE program inside a DESQview window. In the latter case, the CD-ROM is accessible only in that one window, but

that's often good enough.

I prefer to have the WORM drive available in all windows, and its system and BIOS programs take up enough of that high memory area that there's no room for MSCDEX.EXE up there. If I run MSCDEX.EXE to low memory before loading DESQview, the CD-ROM drive is available to all windows, but the largest window I can make is 470K bytes or so, which is not enough.

Eventually, I hope, both Maximum Storage and the various CD-ROM drive vendors will come up with a way to let their access programs load most of themselves into expanded memory *above* 1 MB. Clearly this is the way things must go; there are far too many uses for that magic area between 640K bytes and 1 MB. For now, though, I have to make hard choices.

Anyway, Amdek wanted their drive back, and about then the Bureau of Electronic Publishing arranged for me to get a pair of Denon CD-ROM drives. Unlike the Laserdek drives, which are *almost* SCSI, the Denons are "real" SCSIs; they'll work just fine on the SCSI port on a Mac. Mine came with a PC SCSI board and software for both the PC and the Mac. I've used them with the PC and have also daisy chained them off the Mac's CD-ROM drive; they work fine with both machines.

Installing them in the Z-386/25 was a snap: just drop the board in and run the Install program. Some of the instructions were obviously written by people who speak English as a second (or more likely third) language, but they're clear enough. Once I had them set up in a daisy chain as drives D and E on the Z-386/25, I tested them with Microsoft Bookshelf, the Sherlock Holmes disk, Grolier's Encyclopedia, and several others. I had no

Image Processing for the PC AT or PS/2. Software identical in every respect.

— Fred Molinari, President

QuickCapture gives you a choice of host platforms.

Now you can keep your options open when it comes to choosing a computer platform for image processing.

Because QuickCapture™ on the PC AT® and PS/2™ is strictly software, file, and function compatible.

Just think, applications developed on one machine can be used instantly without modification on the other.

So, add QuickCapture to your choice of platform and you've just added the industry's most versatile frame grabber for all your scientific and commercial applications.

QuickCapture for the PC AT or PS/2. The likeness is uncanny.

	DT2953	DT2855
640x480 resolution	yes	yes
Input/Output LUTs	yes	yes
Input type	RS-170, NTSC	RS-170, NTSC
File formats	TIFF, PCX	TIFF, PCX
Compatible bus	MCA	PC AT

Call (508) 481-3700
In Canada, call (800) 268-0427

FREE 1990 Image Processing Handbook

OVER NIGHT DELIVERY

DATA TRANSLATION®

Circle 72 on Reader Service Card

Listing 1: *The Chaos Manor CONFIG.SYS and AUTOEXEC.BAT files.*

```
CONFIG.SYS
LASTDRIVE=Z
DEVICE=C:\QEMM\QEMM.SYS    RAM
FILES=24
BUFFERS=1
DEVICE=C:\QEMM\LOADHI.SYS /R:1 \BIN\ANSI.SYS
DEVICE=C:\QEMM\LOADHI.SYS  C:\DENON\EXXTF210.SYS
    /D:DENON-CD/N:2/C:10


AUTOEXEC.BAT
PAUSE
VERIFY ON
BREAK=ON
C:\QEMM\LOADHI C:\BIN\POURVOUS.COM
C:\TECMAR\VGA PS/2
PATH=C:\QEMM\;C:\;C:\BIN;C:\NORTON;C:\DV;C:\QW;C:\MOUSE1;
REM  SET CAT=c:\PCLIB\
PROMPT $P $T$H$H$H $G
BE SA WHITE ON BLUE
C:\QEMM\LOADHI /R:2 BUFFERS 15
C:\QEMM\LOADHI /R:1 C:\MOUSE1\MOUSE
REM C:\MOUSE1\SETSPEED /P2 /FC:\MOUSE1\MOUSEPRO.FIL
CHKDSK
C:\QEMM\LOADHI C:\DENON\MSCDEX /D:DENON-CD/M:20
REM SET CDPATH=C:\BKSHLF;D:\SOFTWARE;D:\BOOK
REM DV
```

problems. They're perhaps a bit faster than the Laserdek drives, but it's not a big difference.

They work fine together. I had real trouble daisy chaining Laserdek drives: sometimes they'd work, and sometimes they wouldn't, and I never learned why. The Denons gave no problems at all. Part of my test was to run them in a DESQview window, and that worked fine, too. The only problem I had was that Microsoft Bookshelf is memory-resident and you invoke it by pressing the Left-Shift-Alt keys, and Shift-Alt is understood by DESQview to mean that you're making a keyboard macro. I believe, though, that both Bookshelf and DESQview have ways to reconfigure the hot keys, so it should be simple enough to set them up so they don't conflict.

The Bureau of Electronic Publishing—a private company specializing in CD-ROMs and CD-ROM drives—is currently recommending Denon drives, and Microsoft offered me a Denon to replace the Laserdek. I haven't had these long, but they look good to me. Incidentally, like the Laserdek, the Denons have audio output, and they come with software to let you play compact disks through any stereo or boom box. Works fine, too.

Z-386/25 QEMM-386 Configuration
I have had a number of requests for the exact system I use with DESQview in my Z-386/25. What you see in listing 1 works well with the *new* QEMM.SYS from Quarterdeck. (Note: The DENON line in CONFIG.SYS is in two parts, but it is in fact one long line.)

This works with my current Tecmar VGA board; other boards may require you to exclude certain high memory areas.

QEMM-386 and Manifest
The new QEMM-386 memory manager from Quarterdeck is a delight. It's fast and darned near bulletproof. QEMM-386 will turn all (or as much as you want) your high extended memory—memory above 640K bytes and below 1 MB—into expanded memory usable by any EMS-capable program. It also comes with the LOADHI programs—both .SYS and .EXE—that will put drivers and such up there. Note in my CONFIG.SYS and AUTOEXEC.BAT files in listing 1 that I've packed just about everything into that high memory area.

Quarterdeck also has Manifest, the ultimate 386 system snooper program. Run Manifest, and you'll find out everything you ever wanted to know about your system, such as: Exactly where is the video RAM? Where does DOS put a thing like POURVOUS.COM? Which interrupts do what? And lots more. When I first got Manifest, I spent a couple of hours just poking around with it. It's fun.

QEMM-386 works with nearly everything. It's true that some versions of Microsoft Windows try not to work with QEMM-386, but you can get around that by renaming the program. Instead of QEMM.SYS, rename it to FRED.SYS, or any other .SYS you like, and change the CONFIG.SYS statement to match. QEMM-386 won't mind, and since all

Microsoft is doing is testing for the name "QEMM," Windows/286 won't know what you've done. Windows/386 won't work with QEMM-386 at all, but that's no loss, as I've found Windows/386 to be unreliable.

If you have a 386, you really need QEMM-386 and Manifest, which get a User's Choice Award.

Bible Study

It was a remarkable coincidence. I'd just got the Biblesoft PC Study Bible working on the Z-386/25 when Father Wilcox came to see Roberta. Although it was an unexpected visit, it worked out well: he helped test the program.

The PC Study Bible consists of the Old and New Testaments. The following versions are available: the American Standard Version (ASV), the King James Version (KJV), and the New International Version (NIV). There's a complete concordance of each and a bunch of software to let you look for just about anything you like. You can ask for all instances of the word *poor* in the New Testament, or of *begat* in the Old, or whatever else strikes your fancy. Having got those, you can have the text from two different versions side by side on a split screen.

There's also a word processor with cut-and-paste capabilities and a smaller notepad that you can use anywhere in the program. This is one great aid to scholarly essays or sermons.

The whole thing works smoothly and quickly; on the Z-386/25, it's lightning fast, almost instantaneous. Naturally it's slower on an XT, but it will still get the job done.

The program isn't small. You need 2 MB of storage for the software and one Bible version; each additional version requires another 1.5 MB. You can also get an optional add-on, the Greek-Hebrew Dictionary, with Strong's translation and numbers and The Englishman's Greek-Hebrew Concordance. That adds another 2 MB.

The result is that you have all the tools of scholarship at your fingertips. Father Wilcox was impressed, so much so that I suspect we'll buy him a copy for his birthday.

The interface is surprisingly good, but there are a couple of glitches. For instance, if you tell it to change versions, it will do it almost instantly—but you have to remember those three-letter abbreviations (e.g., KJV) because the program will accept nothing else, and pressing the help key gets you the unhelpful message "That Version Not Available." Still, it's

not all that hard to write these things in a notebook until you've memorized them.

All told, I have had a great deal of fun with this program, and I think anyone with an interest in the Bible, whether as revelation or as literature, would enjoy it. I've uploaded a demonstration version of BIBLE.ARC into my "tojerry/listings" topic on BIX.

AutoCAD

There aren't many well-known microcomputer applications written in Lisp, but Autodesk's AutoCAD is a major exception. AutoCAD is literally the standard that all other CAD programs are compared with. It's an extremely flexible program, with a great number of bells and whistles and nearly any feature you would want, including a programming language.

Of course, AutoCAD is more program than many of us need, and with all its features, it can be a bit formidable when you first try learning it. Fortunately, there's an alternative. AutoSketch is a cut-down version of AutoCAD. It has overlays, but not as many. It has the ability to store parts as objects, so you can bring them back wherever you want. It has dimensions, rotations, springing curves through a set of points, and most of the other things you'd expect.

Best of all, the drawings created with AutoSketch can be read into AutoCAD, if you ever want to go to the full-blown system. You can then get AutoShade, AutoSolid, and suchlike to make as complex a set of drawings as you'd ever want. I'm using AutoSketch to do the preliminaries on a new project. If it gets complex enough, I can change over to AutoCAD; by then, I'll be a lot more used to using CAD systems.

The Fax Dilemma

We have entered the Age of Fax. I think the only person I know who doesn't have a fax system is me, and I'm weakening. I've held out this long because, as Roberta points out, we already get more mail than we can handle; a fax system will generate even more paper. Also, I don't really want to get another telephone line to be tied up by a fax machine.

There's nothing to be done about the paper storm, but the telephone-line objection is a solved problem. The Voice/Data Switch from Rainier Technologies will do the job nicely.

The Voice/Data Switch is a black box about the size of an external modem. You bring your phone line into it and then connect one output to your phone and the

other to a modem or fax machine. You then record a message into the box. When someone calls, the box answers the phone and tells callers to dial 3 to talk to someone or wait for the tone if they want to send a fax (or modem message). That's it. The box switches from one to the other, thus allowing one phone line to serve for both data and voice.

The Voice/Data Switch will work fine with a private-branch-exchange or KSU internal phone system; just plug that in instead of a phone. You can even have another phone connected at the Rainier unit if you like.

This wouldn't be suitable for an office that gets a lot of fax traffic, of course, but for someone who doesn't get much, it will do the job nicely. Naturally, you can send a fax out with this system. While the fax part is in use, anyone who picks up the voice phone will hear a busy signal.

In my case, I'm connecting this to the

new Prometheus ProModem 2400 MFAX fax/modem system for the Mac. I can then use the Mac scanner—either the Apple unit or the new Seikosha scanner—to get an image for transmission, and of course incoming stuff goes into a data file for printing on the LaserWriter. The Prometheus is both a fax and a standard 2400-bps modem, and it works fine. More on it another time.

Winding Down
As usual, I've hardly made a dent in the piles of stuff here. The game of the month is Action Stations (Conflict Analytics, 119-2 Norfolk Rd., Naval Air Station, Alameda, CA 94501). This is a first-class simulation of naval surface warfare as it might have been in, say, 1930, prior to the development of naval aircraft.

The book of the month is actually a pair by Rick Cook: *Wizard's Bane* and

The Wizardry Compiled, from Baen Books. Imagine a typical Silicon Valley nerd transported into a world of magic—and the magic is subject to programming by a skilled hacker. Good fun.

Next month, a whole bunch of stuff, including reports on the American Association for the Advancement of Science meeting. ∎

Jerry Pournelle holds a doctorate in psychology and is a science fiction writer who also earns a comfortable living writing about computers present and future. Jerry welcomes readers' comments and opinions. Send a self-addressed, stamped envelope to Jerry Pournelle, c/o BYTE, One Phoenix Mill Lane, Peterborough, NH 03458. Please put your address on the letter as well as on the envelope. Due to the high volume of letters, Jerry cannot guarantee a personal reply. You can also contact him on BIX as "jerryp."

Pretty.

Mission Critical Workstation 1448: 9 option slots, 2 drive bays and a 95-132 VAC, 47-63 Hz power supply.

A PC that looks good in your office won't look good for long out in the plant.

Where heat murders microprocessors. Dust decimates disk drives. And vibration victimizes video cards.

Which can wreck your entire operation.

That's why Texas Microsystems offers a complete line of reliable, IBM-compatible products and systems specifically engineered for those brutal industrial environments that eat pretty PCs for breakfast.

To ensure maximum reliability we design and manufacture from scratch practically everything that goes into our systems. Like passive backplanes—which, by the way, we pioneered for PCs in 1983. These backplanes accommodate a full complement of convenient, plug-in components. And they're why our mean time to repair is a phenomenal 10 minutes.

You won't find passive backplanes—or lower MTTR—in any of the leading office PCs.

We also build industrial-strength option cards to handle myriad functions. As well as 286, 386™ and 486™ CPU cards in a full range of processor speeds. Our CPU card designs use Very Large Scale Integrated circuits and programmable array logic devices to reduce component counts by 50–60%. Which enhances reliability and resistance to physical stress. And ultimately contributes to our remarkably long mean time between failures: 70,000–

Ultra-fast 32-bit (80386) AT™ equivalent CPU board B386S. Available at 16, 20, 25, 33 MHz.

100,000 hours, calculated against the MIL Standard Handbook 217E.

You won't find that kind of card selection—or MTBF—among the leading PC makers.

You won't find them torturing their systems like we torture ours, either.

Not only do we perform extensive "shake, rattle and roll" tests on each new design, but we pretest all our systems. Burning them in at 55°C for 42 hours straight.

Mission Critical Benchtop 2003: 10 option slots, 2 drive bays and a 95-130 VAC, 47-440 Hz power supply.

Pretty reliable.

Just to make sure they can take the heat at your plant.

What's more, we shock-mount our disk drives to stand up to vibrations surpassing Richter scale proportions. And we use only high-reliability power supplies that can go 100,000 hours, MTBF.

With all that reliability designed into our products, is it any wonder that we guarantee better support than the leading PC makers? Every system we offer comes with a full one-year, on-site warranty. Theirs don't. We also offer a toll-free number for technical and sales information, a regional network of sales engineers, engineering support for systems integration and a guarantee to meet shock specs. And, of the leading PC makers, Texas Microsystems has the longest history of design using Intel microprocessors: 15 years in all. That's why you'll find our systems hard at work in harsh operating environments at 70 of the Fortune 100 companies.

Granted, the leading office PCs may be prettier than ours. But our industrial-strength systems are designed to be more reliable. And that makes our systems look a lot better where it really counts:

Your production line.

Mission Critical Rack-mount 2001: 10 option slots, 3 drive bays and a 95-130 VAC, 47-440 Hz power supply.

It's-No-Comparison Chart			
	Texas Microsystems	COMPAQ	IBM
Passive backplane	YES	NO	NO
100,000-hour-MTBF power supply	YES	NO	NO
Shock-mounted disk drives	YES	NO	NO
Maximized MTBF	YES	NO	NO
Positive pressure, filtration	YES	NO	NO
Operation at 55°C	YES	NO	NO
42 hours burn-in at 55°C	YES	NO	NO
Maximum expansion slots available	14	5	5
1-year, on-site warranty	YES	NO	NO
Toll-free support number	YES	NO	NO
Regional sales support	YES	NO	NO
"Shake, rattle and roll" testing	YES	NO	NO

For technical or sales information, call

1-800-627-8700

Texas Microsystems Incorporated

10618 Rockley Road • Houston, Texas 77099
713-933-8050 • Fax: 713-933-1029

Circle 271 on Reader Service Card

TERMINAL EMULATION FROM TERM ISN'T THE NEXT BEST THING TO BEING THERE. IT IS BEING THERE.

With TERM, you can operate remote systems exactly as if you were there.

TERM is considered the industry standard for exact terminal emulations, but emulation isn't the only issue – regardless how precise it is. The real objective is access, and accessibility requires that your machine act like their machine.

Exact Emulation.

With SCO Color Console and TERM, anything that runs on Xenix can now be operated remotely from your PC, Macintosh or Unix box. TERM literally puts you in front of the remote system. You can now take advantage of the full range of keyboard functions and screen attributes including function key programmability and complete color support. In short, your terminal becomes their terminal. This is access.

In addition to SCO ANSI, TERM features very exact VT220, VT102, Wyse50 and nine other terminal emulations on any physical display device. Other emulation features include a fully remappable keyboard, 132 columns, transparent print, graphics line draw and multinational character support.

Network & Asynchronous Access.

Exact emulations are merely one feature in a communications program which is considered the industry standard for asynchronous communications. And now Century Software introduces TermNET- all of the features of TERM in a network version.

TermNET operates identically to TERM and provides direct network connectivity featuring multisession TCP/IP, OSI and X.25 support. Both TERM and TermNET provide intelligent file transfer for fast data movement and a powerful, customizable script language that automates every aspect of communication sessions. They run identically on machines made by over 100 manufacturers, providing keystroke-for-keystroke compatibility across multiple operating systems including: MSDOS, UNIX, Xenix, Macintosh, VMS, BTOS, MS-Windows and X Windows.
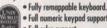
Access & Connectivity Solutions.

Century Software is dedicated to meeting the challenges of systems interoperability. With the addition of TermACS (Asynchronous Communications Server) and TERM for X Windows to our family of software products, we are creating *The Access Standard* – a complete line of communications tools that provide state-of-the-art asynchronous and network connectivity. This means access to information from any system, any network.

CENTURY IS THE ACCESS STANDARD.

TERM AT A GLANCE
- Multi-session Network & Async Access.
- Direct TCP/IP, OSI & X.25 network support.
- Exact VT220, VT102, Wyse50 and SCO ANSI on all systems.
- Fully remappable keyboard.
- Full numeric keypad support.
- Full color support.
- Automatically restartable file transfers.
- State-of-the-art Lempel-Ziv-Welch data compression.
- 38.4K sliding window file transfers.
- KERMIT protocol for mainframes.
- XMODEM and YMODEM protocols for bulletin boards.
- Wildcard file send/receive capability.
- Auto-login, dial/redial modem control.
- Powerful script language for customized applications.
- Performs unattended polling/file transfers.
- Remote maintenance capability.
- Electronic mail/TELEX/FAX Easylink/MCI Gateway.
- Superior X Windows xterm Replacement.

CENTURY SOFTWARE
The Access Standard

5284 South 320 West,
Suite C134
Salt Lake City, Utah 8410
Phone: 801-268-3088
Facsimile: 801-268-2772

UNITED KINGDOM	FRANCE	FRANCE	BELGIUM	NETHERLANDS	ITALY	WEST GERMANY	AUSTRALIA
Systems Marketing	Tauris Data	Top-Log	Top-Log	T-Point Software	ESA Computers	Tauris Data	Qunix
Phone: 0488-681004	Phone: (1) 30-21-55-05	Phone: (1) 42-04-21-18	Phone: 02/872-22-40	Phone: 074-912470	Phone: 0541-741113	Phone: (6121) 509392	Phone: (07) 831-86

Circle 47 on Reader Service Card (DEALERS: 48)

PROWLING THE NETWORKS

Getting on the
Usenet network, and
what to do when
you've gotten there

U senet, that anarchistic collection of computer sites that nobody seems able to explain adequately, is the network most information-starved Unix system owners aspire to connect to. Once your system is on Usenet, you can read and post news, gather the free software posted to comp-.sources and alt.sources, keep your network maps up to date with comp-.mail.maps, and exchange E-mail with people from around the world.

Getting on the Network
To get on Usenet, the age-old advice is, "Ask around." Unix users with network addresses are generally the best source of information about where to find a nearby mail or news feed (a site that agrees to supply you with a network connection). Computer science centers at local universities are often good places to start.

Many sites are set up for public-access mail and Usenet gateways, and these might help if you can't find any local feeds. There are also sites that allow dial-up connections to their Usenet-connected Unix systems—a useful idea for people with personal computers who don't want to run Unix just for Usenet purposes. A list of many of these systems (some of which request donations) is available by setting up a UUCP connection to the Telebit-equipped system lgnp1 at (215) 348-9727, using a log-in name of nuucp and no password, and requesting the file /usr/spool/uucppublic/nixpub via UUCP. Note that no public BBS or regular dial-up log-in is available here. This is only for UUCP connections, courtesy of the donated

time and equipment of Phil Eschallier.

A good node to hook up to might be UUNET (3110 Fairview Park Dr., Suite 570-B, Falls Church, VA 22042, (703) 876-5050, fax (703) 876-5059, or info@uunet.uu.net), a very large and specialized site with connections not only to Usenet, but also to the Internet and other major networks. UUNET is run by a nonprofit organization originally sponsored by the Usenix Association and exists solely as a communications link. It provides mail and news feeds, and it has over 600 megabytes of public domain software available for direct access or on tape. It also runs a commercial TCP/IP network service. Subscribers can get substantial discounts on certain modems and books, and a special rate for low-volume users was recently instituted. Dozens of Telebit modems, a high-speed T1 link and X.25 connections, and over 2.5 gigabytes of disk space make UUNET a unique system.

Getting the News
Getting the software needed to run a Usenet node seems like a classic dragon-and-egg problem, since it's distributed via Usenet. But once you have a simple UUCP connection to a Usenet site, they can just send you the source code, since the NetNews software itself is public domain. Xenix users have it a bit easier; they can get a precompiled version of the NetNews software directly from The Santa Cruz Operation.

Getting NetNews running is complex but not arduous; it's certainly more than can be covered in this limited space. I refer readers in search of more detail to Ben Smith's article "The Unix Connection" (May 1989 BYTE) and to the books *Managing UUCP and Usenet* by Tim O'Reilly and Grace Todino and *Using UUCP and Usenet* by Grace Todino, both published by O'Reilly and Associates, 632 Petaluma Ave., Sebastopol,

continued

CA 95472, (800) 338-NUTS, (707) 829-0515, or ora!nuts. Once you've successfully exchanged news messages with your news feed, all you have to do is post your site data to the group news.new-sites, with a copy to rutgers!uucp-map, and you're on the network.

So, what do you do once you are able to receive files from the network? The first thing you might do is ask your news feeds if they have the source code to a number of important communications-oriented programs, without which you will be working a lot harder than necessary (and if they don't have it, you can post a message in comp.sources.wanted).

Free and Useful

Written by Peter Honeyman and Steve Bellovin, pathalias computes relative mail paths from your computer to all others known on the network. It works by parsing the map data posted to the group comp.mail.maps and creates a database usable by other programs, including uuhosts, smail, and elm. The completed database contains not only what are considered to be the quickest paths, but also their relative cost. With over 18,000 entries in the database and all the possible permutations, you'll probably need at least 4 MB of memory to run it without excessive paging or swapping.

Created by John Quarterman, uuhosts is a large shell program that helps manage the information in the comp.mail.maps data files. It creates an index of all sites on the network and lets you look up all available data about a site on the fly. Once you've tried perusing the data with grep, you'll know why Quarterman wrote this program.

Many systems have antiquated mail delivery programs that can't properly handle network addresses and security concerns. The smail program solves these problems fairly well, although you might have to be creative when installing it on some systems. The program is configurable as issued, and, of course, you always have the source code to look at. The current released version is by Christopher Seiwald, as far as I can determine.

Originally written by Dave Taylor, elm (for Electronic Mailer) has generated so much interest that Hewlett-Packard (for whom it was written) has turned it over to a group called the Usenet Community Trust for further development. It's a screen-oriented mail program that you can use with AT&T Mail Forms as well as with regular Unix mail. It can use the pathalias database for outgoing mail and can reverse the path of incoming mail to help ensure that replies will get to the original sender. There are zillions of options that let you customize elm's performance and even its screen appearance. The program can even file and answer mail for you automatically.
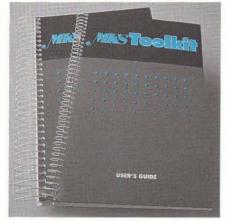
Another program that I have found very useful is xcomm, by Larry Gensch (lar3ry on BIX). It is a Unix-oriented telecommunications program, mainly oriented toward calling non-Unix systems. So that logging on can be done automatically, xcomm keeps a list of phone numbers and start-up scripts, which is similar to the way that UUCP works. It has cu-compatible put and take commands and lets you upload and download both binary and text files, using XMODEM and CIS protocols. You can start or end disk logs of interesting material while you are on-line. Its script language allows you to automate sessions. The program fills in the gaps left by standard Unix commands and lets you communicate with any system that supports the ASCII character set.

Other Networks

Usenet is certainly not the only Unix-oriented network, although it is the largest. Some other (but still by no means all) important networks include the Internet (which itself comprises ARPANET, CSNET, MILNET, and NSFNET), BIT-NET, CompuServe, FidoNet, and MCI Mail. Geographically oriented networks include ACSnet (Australia), EUnet (Europe), IUnet (Italy), JUNET (Japan), and UKnet (United Kingdom). While most of these are Unix-oriented, quite a few are not, but all can be reached via UUCP or TCP/IP connections from any Unix system. The trick is to have the appropriate mail software and to know how to reach the "gateway" machine to the target network while using the correct address.

That brings up another problem. Unix users using UUCP are accustomed to addresses like gecko!lizard!serpent!reptile!dragon, where you must specify an explicit path to ensure mail delivery. The preferred addressing method for the 1990s is called "domain addressing," which leaves the actual delivery mechanism out of the address. Thus, while my computer has been known for years as simply "infopro" under UUCP-Net conventions, you can now address it as infopro.uucp, signifying that it is the infopro machine in the UUCP domain. Notice that the "bangs" (! characters) are gone under this scheme. My new, improved, full domain address, therefore, is now david@infopro.uucp, in place of

the older `infopro!david`.

There are many other domains, each of which is a high-level gateway that knows how to route messages. For instance, you could have an address such as `foo@bar.klunker.com`, which would signify that the destination is a person with the log-in name foo on the computer called bar at a company named klunker (the suffix `.com` means it's a company or a commercial entity of some kind). Or you could send mail to `joe@finagle.smart-college.edu` (joe attends smart-college, an EDUcational institution, and has an account on a machine called finagle), or to `singh@curry.calcutta.in` (a person named singh at a computer called curry in Calcutta, India).

In each case, your software only has to know how to translate the domain (`.com`, `.edu`, or `.in`) into a machine destination. From that gateway machine, your message will be routed (internally to the domain) to the destination machine. In many cases, the information provided in the maps will go even further, allowing your computer to generate a path directly to `klunker.com`, `smart-college.edu`, or `calcutta.in`.

For more information on the wonderful and confusing world of network addressing and gateways (the above is only a small sample of the possible confusion), a recently released book by Donnalyn Frey and Rick Adams called *!%@:: A Directory of Electronic Mail Addressing and Networks* (O'Reilly and Associates, 1990) is a must. The book lists full details on over 100 different networks and has maps, contact, and gateway information, as well as many pages on all the addressing schemes. Rick Adams is the president and technical director of UUNET Communications Services and is in a good position to know the tricks for getting around the networks.

Over the next few months, I'll expand the discussion of public domain software past communications. I'll talk about the legendary Free Software Foundation and GNU Project, available archives and where to find them, and the work and benefits you can get from using "freely available" software. ■

David Fiedler is publisher of the Unix Video Quarterly *and the journal* Root, *as well as coauthor of the book* Unix System Administration. *He can be reached on BIX as "fiedler."*

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

Circle 247 on Reader Service Card

DO-IT-YOURSELF CD-ROMs

If the information you need to access isn't on a CD-ROM, maybe you need to make your own

L ast month I talked about how useful CD-ROMs are if you need rapid access to a vast quantity of information from commercially available sources. But what do you do if the data you need to access is your own proprietary information? You learn to "roll your own" CD-ROM.
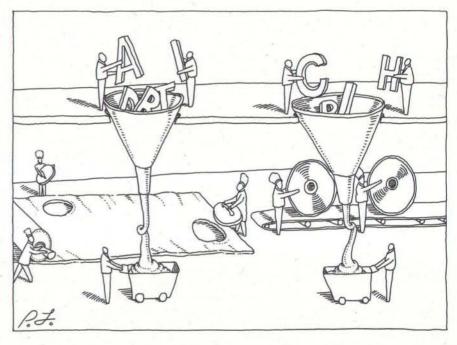
A government agency had this kind of a problem. It simply had too many repair parts to keep track of. As a result, the agency's employees were frequently unable to find the correct replacement parts even though they were available somewhere in the warehouses and storerooms. Managers found that they were often buying parts at retail that were already in inventory because no one had been able to locate them.

Determining the cause of the problem was easy. The agency was responsible for keeping track of hundreds of thousands of parts, which it dutifully listed in catalogs that were printed on paper or stored on microfiche. There were so many entries in the catalogs that the employees couldn't be sure whether they were really finding the proper part or whether they were missing an item that was on a change sheet or a missing page.

The paper catalogs and parts listings were the biggest problem. A search of the piles of paper was doomed to failure. Even the handling and storage of the paper that was required for the catalogs was a problem. A solution was needed.

ROM to the Rescue

Although the agency had thousands of items to keep track of, the inventory was fairly stable. New items were added to the list slowly, and old items were removed only occasionally. What was needed was a compact, easy-to-use way to find the repair parts. A CD-ROM was an obvious choice.

Moving the parts catalogs to CD-ROM allowed the employees to look for a part by its stock number, its name, or by the name or stock number of the item it was installed in. The CD-ROM software allowed users to read descriptions of the item, see ordering instructions, and even see a drawing of the item if necessary. Suddenly, all that was required to serve as a catalog was a personal computer with a CD-ROM drive and a user with minimal training.

Doing It Yourself

While the government agency had the help of a consultant to convert the parts data for use with a CD-ROM, such conversions are rapidly becoming easy enough that some companies can do the work themselves. At the same time, the cost is dropping so fast and so drastically that publishing on a CD-ROM may become cheaper than publishing a catalog on paper. It all depends on the application.

This is not to say that the cost is insignificant. A company that has to convert records on paper to digital information will have a great deal of work to do. The information may have to be keyed in by hand, one record at a time, which can be expensive.

Of course, many businesses already have their information digitized. It may be in the form of an inventory, for example, which can easily be stripped of unnecessary fields and used as the information for the CD-ROM.

Once you've done this, the next important step is deciding what your users are going to do with the data. What kind of searches are you going to do? What will

continued

ITEMS DISCUSSED

Prices are not included due to the individualized nature of the data conversion, data preparation, and premastering processes.

FULL-SERVICE CD-ROM DEVELOPMENT:

Network Technology Corp.
7401-F Fullerton Rd.
Springfield, VA 22153
(703) 866-9000
Inquiry 1101.

CD-ROM DATABASE SOFTWARE:

Graphic Knowledge Retrieval System
Text Knowledge Retrieval System
KnowledgeSet Corp.
888 Villa St., Suite 500
Mountain View, CA 94041
(415) 968-9888
Inquiry 1102.

MediaBase
Crowninshield Software, Inc.
1105 Commonwealth Ave.
Boston, MA 02215
(617) 787-8830
Inquiry 1103.

Optical Build
Optical Convert
Optical HSG-Build
Optical Prep
Optical Publishing, Inc.
155 West Harvard St.
Fort Collins, CO 80525
(800) 869-7989
(303) 226-3466
Inquiry 1104.

the search criteria be? This information is critical in deciding how the data will be organized and indexed, and how useful the resultant CD-ROM will be.

Searches and Indexes
Deciding on the search criteria probably isn't as hard as it may seem at first glance. Essentially, the process consists of deciding what the users will be looking for and where they might find it. A parts listing, for example, might be set up so that you can search for the part number, the name, the revision or modification number, and the part number of the parent assembly. Other types of data require other search methods. You may even want to do a full text search in some cases.

The details of how you will approach this process depend to some extent on the abilities of the database software you plan to use. Some of the database packages designed for use with CD-ROMs let you search for two words in the vicinity of each other, some allow logical operators, and some only allow simple text searches.

Once you have decided what you want to look for and how you want to find it, you need to get your data into shape. It has to be in a format that the company that makes the CD-ROMs can use. Normally, if the data is in a flat ASCII file, a dBASE file, or a comma-delimited file, you needn't do anything. Your data is already in an acceptable format. Once you have chosen how you want to present

your data, it's time to get it ready for the CD-ROM company.

The Four Steps
I spoke to Tom Vreeland, president of Network Technology in Springfield, Virginia. Network Technology provides the Lasertex software that many CD-ROM developers use to support searches. It also provides a complete CD-ROM development package, from preparing the data to providing the finished product.

According to Vreeland, there are four essential steps to producing a CD-ROM for distribution: preparing the data, premastering, mastering, and replication. Any company with the proper equipment can perform the first two steps and save the cost of paying an outside firm to do it. Mastering and replication involve the production of the actual disk, which requires specialized equipment. This must be done by a compact disk production company.

Once your company data is in electronic format (either it was already on disk or you had it keyed in), you must prepare it for premastering. This means that you have to take the information you developed about search criteria and indexing and put it into use.

Companies such as Network Technology, Crowninshield Software (Boston, MA), KnowledgeSet (Mountain View, CA), and Optical Publishing (Fort Collins, CO) provide database software and development tools that allow companies

We'll take your stats and make you the most valuable player in your league.

to develop their own CD-ROM search routines using packaged software. These packages provide an easy-to-use interface with menus and windows, and they usually support the display of graphics images on computers equipped to display them.

Once the data-preparation step, which includes testing of the search software, is complete, the next step is premastering. This process involves adding the indexes and the error-correction information that will be needed to support CD-ROM use and transferring the information to either a nine-track tape or a WORM (write once, read many times) disk.

Most software will support CD-ROM development on an IBM AT–compatible computer. A few packages will let you use a minicomputer or a mainframe, which will speed up the process. In either case, the transfer medium must hold all the data, the indexes, and the error-correction information, which can amount to as much as 800 megabytes for one CD-ROM. This is why you need tape or a WORM disk.

Once you have delivered the tape to the CD-ROM production company, the pro-cess is out of your hands. The company will create the master used for manufacturing the disk itself and then injection-mold the distribution copies. You will receive a shipment of the finished product, ready for sending out to your customers.

The Costs

Depending on how much of the work you can do yourself, and on what shape your data is already in, CD-ROM production can be relatively inexpensive. According to Vreeland, the entire process, including production of the finished disks, can cost less than $10,000. The exact costs are difficult to determine, because they depend on the amount of data, the work to be done, and the royalties on the database software.

Network Technology charges $13 per disk—a price that covers disk production and the royalties. Other suppliers have other ways of charging, so exact comparisons are difficult unless you have detailed information on your specific circumstances.

One often-overlooked cost is the addition of a CD-ROM drive to the computers that you'll use to access the CD-ROM you are developing. In the past, these drives were expensive, but the price has been dropping. Vreeland pointed out that CD-ROM suppliers, including his company, are trying to find ways to increase sales of the CD-ROM drives so that more people can use the CDs they produce.

Because of much lower prices, it's much easier to convince your customers to buy a CD-ROM drive to use your catalog. After all, they don't want to be inundated by catalogs any more than you do. And they do want to be able to find products and inventory easily using a minimum of staff time. CD-ROMs can make that happen. ■
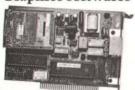
Wayne Rash Jr. is a contributing editor for BYTE and a member of the professional staff of American Management Systems, Inc. (Arlington, VA). He consults with the federal government on microcomputers and communications. You can contact him on BIX as "waynerash," or in the to.wayne conference.

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

When your memory's low,

we've got the chips

PC Cultivation.

Silicon salad days.
(Or, how we mind our Peas & Cukes.)

Way up north in the fertile crescent of Marlow, NH (pop. 562), we know how to dig down deep. Which is pretty difficult (even for our celebrated 20 mole team) since the bedrock's just inches below the surface and the growing season's shorter than the day is long. But, with the winds of change blowing non-stop through the micro fields, you can't just scatter your seed any which way. You have to put down roots! We ought to know—we were the first company to sell peripherals and software exclusively for IBM personal computers. So when customers call us for product specs, prices, or technical assistance, they're dealing with a company that has its paws (and peas) planted firmly in the ground.

A window-based system you can grow with.

Up here in Marlow, our imagination isn't the only thing that's fertile. In fact, it's small potatoes compared to our PC Connection Mint Garden which comes complete with soil and seeds for growing a luscious crop of Spearmint, Peppermint and Lemon Balm right on your favorite window. Put a fresh sprig in your favorite beverage and we're sure it will add a delightful *je ne sais quoi* (which means that we have no idea what it will add—but it'll taste good). It's free to everyone who places an order of $500 or more between now and June 30.
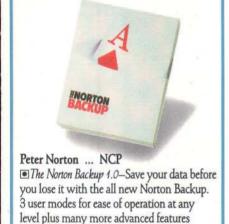
Grow your own in a PC Connection Mint Garden.
Offer not available outside the Continental U.S. or
to accounts on net terms. One per customer.

PC CONNECTION®

to make your PC go.

IOMEGA ... 1 year

5116	Bernoulli II Single 44 Meg Internal	$995.
5117	Bernoulli II Dual 44 Meg External	1969.
5113	44 Meg Cartridge Tripak (5¼") ...	249.
2499	PC2 Card (controller required) ...	169.

Mountain Computer ... 1 year

2917	40-60 Meg Internal Tape Drive....	379.
5502	83-152M Ext. Tape Drive........	759.
5500	83-152M Int. Tape Drive........	629.
6153	DC2120 Cartridge (5 pack).......	135.
5190	DC2000 Pre-formatted Cartridges ea.	35.

Pacific Rim ... 1 year

| 5010 | 1.2 Meg External (for PS/2's)... | 215. |
| 6602 | 1.44 External (for PC/XT/AT)... | 239. |

Plus Development ... 2 years

3105	Hardcard 20 Meg (49 ms).......	519.
3106	Hardcard 40 Meg (28 ms).......	599.
6425	Hardcard II 40 Meg (19 ms).....	599.
6424	Hardcard II 80 Meg (19 ms).....	699.

Seagate ... 1 year
*FREE PCTV® Hard Drive Installation
Tape with purchase of 20, 30 or 40 Meg
Seagate drive for the IBM PC (not for
AT). Beta or VHS.*

| 2285 | 20 Meg Internal Hard Drive ST225 (w/controller and cables, 65 ms).. | 275. |
| 2286 | 30 Meg Internal Hard Drive ST238 (w/controller and cables, 65 ms).. | 289. |

Peter Norton ... NCP
◉ *The Norton Backup 1.0*—Save your data before you lose it with the all new Norton Backup. 3 user modes for ease of operation at any level plus many more advanced features make your backups easy .... $89.

| 4554 | 40 Meg Int. HD ST251-1 (28 ms).. | 359. |
| 2287 | 40 Meg Int. HD for PC ST251-1 (w/controller and cables, 28 ms).. | 419. |

TEAC ... 1 year

4950	PC, XT 360k Drive (5¼")..........	79.
4951	720k Drive (specify XT or AT, 3½").	79.
4670	1.44 Meg Drive for XT (3½")......	99.
4326	1.44 Meg Drive for AT (includes Bastech software utilities, 3½" copy prot.).	119.

MISCELLANEOUS

Checkfree
| 6360 | CheckFree Xpress.............. | 25. |

CompuServe
| 1676 | CompuServe Information Service.. | 24. |

Cables ... lifetime

1019	Smartmodem-to-AT cable (10 feet)	$15.
5511	Right Angle Printer cable (6 feet)...	15.
1050	Parallel Printer cable (15 feet).....	19.

DISKS

Maxell ... lifetime
2789	5¼" MD2-D 360k Disks (Qty. 10)....	13.
2790	5¼" MD2-HD 1.2Mb Disks (Qty. 10)..	19.
2792	3½" DS/DD 720k Diskettes (Qty. 10)..	15.
2793	3½" DS/HD 1.44Mb Diskettes (Qty. 10)	29.

Intel ... 5 years
The newest modems from Intel provide MNP/ Level 5 support for faster, more reliable communications. Also operates in Hayes-compatible mode for standard communication tasks. *2400EX MNP* $229. *2400 Internal MNP* $199.

Kensington Microware ... 1 year
Master Piece Plus—An elegant power control center with complete surge protection and noise filtering. Includes five outlets, unique low voltage indicator, modem protection, anti-static pad, and swivel base ..... $109.
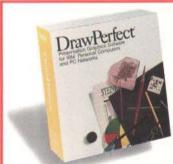
Sony ... lifetime
3291	5¼" DS/DD 360k Disks (Qty. 10)....	12.
3292	5¼" DS/HD 1.2Mb Disks (Qty. 10)...	19.
3297	3½" DS/DD 720k Diskettes (Qty. 10)..	14.
3298	3½" DS/HD 1.44Mb Diskettes (Qty. 10)	29.
6659	QD 2000 Tape Cartridge.........	19.
6712	QD 600A Tape Cartridge.........	27.
6715	QD 6150 Tape Cartridge.........	27.

MEMORY

6556	256k DRAMs (100 nanosecond)...	call
3248	256k DRAMs (120 nanosecond)...	call
4366	1 Meg x 9 SIMMs (100 nanosecond)	call
5510	1 Meg x 9 SIMMs (80 nanosecond)	call
5746	1 Meg Chips (80 nanosecond)....	call

OUR POLICY

- We accept VISA and MASTERCARD only.
- No surcharge added for credit card orders.
- Your card is not charged until we ship.
- If we must ship a partial order, we never charge freight on the shipment(s) that complete the order (in the U.S.).
- No sales tax.
- All U.S. shipments insured; no additional charge.
- APO/FPO orders shipped 1st Class Mail.
- International orders U.S. $250 minimum.
- Upon receipt and approval, personal and company checks clear the same day for immediate shipment of your order.
- COD max. $1000. Cash, cashier's check, or money order.
- 120 day limited warranty on all products. *
- To order, call us Monday through Friday 9:00 AM to 1:00 AM, or Saturday 9:00 AM to 5:30 PM. You can call our business offices at 603/446-3383 Monday through Friday 9:00 AM to 5:30 PM.

WordPerfect ... NCP
◉ *DrawPerfect 1.0*—Comprehensive presentation graphics program from the makers of WordPerfect. Create special effects, self-running presentations, & more. Completely compatible with WordPerfect & PlanPerfect ...... $279.

SHIPPING

Note: Accounts on net terms pay actual shipping.
Continental US:
- For heavy hardware items such as printers, monitors, Bernoulli Boxes, etc. pay actual charges. Call for UPS 2nd-Day & Next-Day-Air.
- For all other items, add $3 per order to cover UPS Shipping. For such items, we automatically use UPS 2nd-Day-Air at no extra charge if you are more than 2 days from us by UPS ground.
Hawaii:
- For monitors, printers, Bernoulli Boxes, computers, hard drives, and power backups, actual UPS Blue charge will be added. For all other items, add $3 per order.
Alaska and outside Continental US:
- Call 603/446-7721 for information.

PC CONNECTION®

MICROSOFT'S NETWORK HEAVYWEIGHT

LAN Manager finally has the right stuff to challenge Novell's NetWare
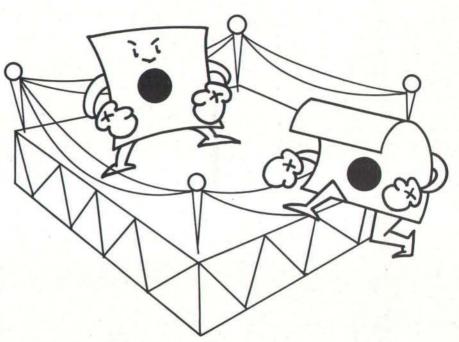
N ovell's NetWare presently holds the LAN title. Microsoft's LAN Manager has aspired to the championship for some time, but its current version, 1.0, has never been a serious contender. LAN Manager 2.0, however, just may have what it takes to make a serious bid.

Microsoft needed to release a radically improved version of LAN Manager, because version 1.0 was scarcely more than a basic LAN operating system. Sure, it let you share files and printers. It even had the major underlying functions that such applications as the SQL Server required.

What it lacked were the many features that would turn it into something that you could trust with large applications and their data. In an all-too-rare moment of vendor honesty, Microsoft even criticized its own product in a "report card" that the firm distributed in its Fall 1989 seminar entitled "Back to School with LAN Manager." While that report card gave LAN Manager 1.0 A's and B's for basic LAN services, the product didn't get above a C+ in such significant areas as administration, security, and fault tolerance.

With all these deficiencies, you might at least expect to get a small system. No such luck; LAN Manager 1.0 was a memory hog that consumed about 150K bytes on each client machine.

LAN Manager 1.0 had another problem that the report card didn't mention: The main commercially available versions, 3Com's 3+Open and IBM's LAN Server, were just too different. Even the

perennial partners Microsoft and IBM bickered over whether such technical goodies as named pipes would be in their respective versions.

Meanwhile, Novell was happily selling its NetWare line. That line wasn't perfectly compatible (see "NetWare's Missing Links," February BYTE), but it was a lot closer than the LAN Manager versions.

Novell even went a step further and shipped NetWare 386, which can take advantage of the power of the 386 processor. LAN Manager 1.0, on the other hand, shared a basic trait with its underlying operating system, OS/2: Both were aimed at the 286 and made no special use of the 386.

LAN Manager 1.0 was also stuck with the slow OS/2 file system, which paid in speed for its DOS compatibility. Perhaps even more damaging is the fact that OS/2 still has not taken off the way that Microsoft hoped it would.

New Strengths
Now along comes LAN Manager 2.0. It isn't perfect, but it addresses all the problems of version 1.0. We like it enough that we plan to install and test the beta version on our lab's biggest server, the dual-processor Compaq Systempro.

Version 2.0's improvements start with those areas that earned bad grades in version 1.0. Its security options, for example, now include such important functions as password aging and the ability to restrict when and where LAN users can log in.

The new version's fault-tolerant features are now competitive with those of SFT NetWare. It can work with uninterruptible power supplies. It includes such data redundancy options as disk and controller mirroring, in which one disk/controller pair can be a constantly up-to-date copy of another.

LAN Manager 2.0 also makes it easier
continued

to manage large LANs, chiefly by offering domains. A domain is a group of servers that share a single database of user accounts. LAN Manager stores that account database on a single server, the domain controller. All log-in requests go through the domain controller; after logging in, users can work with any server they like. There are also now five levels of operator privileges.

Smaller and Faster
While it was adding all these new features, Microsoft didn't forget the product's core. LAN Manager 2.0's client code has slimmed down to under 100K bytes. If you have EMS 4.0 expanded memory, you can even get the client portion down to 60K bytes.

LAN Manager 2.0 is faster than the previous version, too, although here the big gains are in the server code. One of the most significant improvements is LAN Manager's ability to use OS/2 1.2's new High Performance File System. The HPFS, with its B-tree sorted directories, is faster than the DOS file allocation table file system of earlier OS/2 versions. The HPFS also does a better job than the

FAT system of keeping files contiguous, so that such disk improvement techniques as sector prefetching and caching yield better results.

Microsoft has also made a special 386 server version of 2.0 that should yield further performance gains. This 386 version actually makes some changes to OS/2 1.2, which still doesn't do anything special to use a 386.

The main change has to do with where in the processor hierarchy of the 386 (or i486) LAN Manager 2.0 runs. The 286, 386, and i486 offer four levels, or rings, of processes. The most privileged level, ring 0, is where the guts of OS/2, its kernel, run. LAN Manager normally runs in the least privileged level, ring 3, along with all the other OS/2 application programs. This design forces LAN Manager to run under the normal OS/2 security controls and doesn't require any changes to the OS/2 kernel—LAN Manager is just another application.

This design is also potentially slow. For one thing, communication between processes in different rings is, in processor terms, fairly expensive. Also, every time LAN Manager needs any file ser-

vice, it has to go through OS/2.

LAN Manager 2.0 changes all that when it runs on 386 or 486 servers. Its most critical file access code goes into a subsystem in ring 0. This subsystem includes a special 386 version of HPFS, the core of the Server Message Block protocol, and the LAN adapter device drivers. These components bypass OS/2 entirely and manage their own interactions and handle all LAN Manager 2.0 file access requests. (All other requests go to the normal LAN Manager code, which is still in ring 3.) Microsoft further increased the speed of these components by rewriting them to take advantage of 386 instructions.

Even that's not the whole story, because you can move the entire new 386 portion of LAN Manager 2.0 onto a second processor in a dual-processor server. That processor will handle nothing but the well-greased file requests. The primary processor will run the rest of the world: OS/2, the ring 3 part of LAN Manager, and any LAN applications on your server.

This processor arrangement is called
continued

asymmetric because the two processors are definitely not equal. Not surprisingly, many vendors, including Compaq, have recently announced asymmetric, dual-processor 386/486 servers.

A Salute to Consistency
Even with these technical improvements, LAN Manager 2.0 still has a potential multiple-version problem. While its biggest commercial versions remain IBM's LAN Server 1.2 and 3Com's 3+Open

2.0, such vendors as Compaq, Torus, Tapestry, and Ungermann-Bass will also join the LAN Manager 2.0 field.

Unlike the case with their version 1.0 products, the vendors of version 2.0 seem to be trying to stick to a single core. We've talked to a lot of them, and it really does sound like they all plan to stay close to the Microsoft base—when both IBM and 3Com say that they want the LAN Manager core to be consistent, we have to take note.

But that is not to say that these different vendors will offer identical products—far from it. They'll just make their marks by supplying different products and services built around the 2.0 core. 3Com will add such features as mail, internetworking, and support for the Apple File Protocol. We expect IBM to leverage its integration with its databases and mainframes by offering such services as Unified Profile Management, in which the database and the file servers will share user IDs and passwords.

The different versions also will not have completely compatible cores. IBM's LAN Server 1.2 will still have some differences in its application programming interfaces, although we believe that the vendor does plan to make good on its commitment to bring LAN Server and LAN Manager together.

The Byte Budget
All of us who use LAN Manager 2.0 will also pay something for its many new features: memory. The vendors claim minimum requirements of 4 to 6 megabytes; we suspect that 8 MB will be the norm.

We also have to wonder when we'll finally see shrink-wrapped versions of version 2.0. IBM and Compaq claim that they will ship the product in the first quarter of this year, so it should have arrived before this column sees print. We pressed IBM on this point, and the vendor held firm to its March 30 ship date. 3Com will follow in the third quarter of this year; it needs the additional time because of all the extra features that its version will include. Those dates seem believable, if a little optimistic, as we write this column. Microsoft has the software in beta testing, and OEMs have beta versions in their hands.

A Call to Arms
If LAN Manager 2.0 does everything that its press releases say it will, it will have taken a giant step forward. We don't expect it to be the end of NetWare by any means, but LAN Manager 2.0 is impressive. It will be interesting to watch its effect on the LAN marketplace and on NetWare. We'll keep you posted. ■

Bill Catchings and Mark L. Van Name are BYTE contributing editors. Both are also independent computer consultants and freelance writers based in Raleigh, North Carolina. You can reach them on BIX as "wbc3" and "mvanname," respectively.

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

Our Printer Sharing Unit Does Networking!

An Integrated Solution

Take our **Master Switch**™, a sophisticated sharing device, combine it with **MasterNet**™ networking software for PCs, and you've got an integrated solution for printer and plotter sharing, file transfer, electronic mail, and a lot more. Of course you can also share modems, minis, and mainframes or access the network remotely. Installation and operation is very simple.

Versatile

Or you can use the Master Switch to link any computer or peripheral with a serial or parallel interface. The switch accepts over 20 commands for controlling the flow of data. It may be operated automatically, by command, or with interactive menus. Its buffer is expandable to one megabyte and holds up to 64 simultaneous jobs. The

MasterLink™ utility diskette for PCs comes with every unit and unleashes the power of the switch with its memory-resident access to the commands and menus.

Other Products

We have a full line of connectivity solutions. If you just want printer sharing, we've got

it. We also have automatic switches, code-activated switches, buffers, converters, cables, protocol converters, multiplexers, line drivers, and other products.

Commitment to Excellence

At Rose Electronics, we're not satisfied until you're satisfied. That's why we have thousands of customers around the world including large, medium, and small businesses, factories, stores, educational institutions, and Federal, state, and local governments. We back our products with full technical support, a one-year warranty, and a thirty-day money-back guarantee.

ROSE ELECTRONICS

Give a Rose to your computer

Call now for literature or more information.
(800) 333-9343

THE FRUITS OF CONNECTIVITY

It's not enough
to connect computers;
they must also share
information
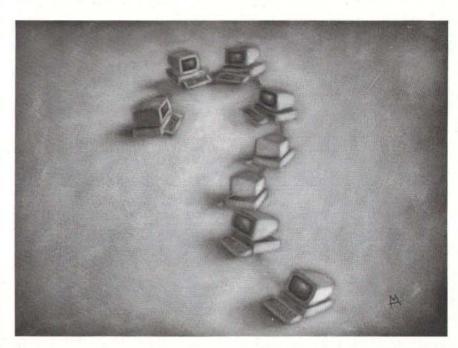
onnectivity. Apple sometimes chants this word like a mantra. It's as if connectivity were the be-all and end-all of computing. With the multiplatform shops of the nineties, connectivity can often make the difference between everyone working together or in isolation.

However, this often overlooks the reason for these interconnections. The fact that the Macintosh now uses LocalTalk, EtherTalk, and TokenTalk misses the real point. Those environments are just the nuts and bolts that tie dissimilar computers together. The real reason for connecting computers is simple: sharing information.

Apple has done a decent job of creating—or encouraging others to create—all those different networking technologies for the Mac. However, Apple should go further. Where is its vision of the new computing made possible by these connections? Where is Apple's white paper on how connectivity makes all kinds of new uses for technology possible?

Telling us that the Mac now works fully in all the major microcomputer, minicomputer, and mainframe networking environments isn't adequate. Apple should also be telling us how these interconnections, combined with the Mac's own ease-of-use benefits, will make new things possible.

Instead, Apple's Knowledge Navigator picture of the future is vague, almost fanciful. Come on, now: What kind of information sharing should all these hookups allow? Do you see a world dominated by new categories of groupware

applications (à la Lotus Notes)? Do you see transparent computing across platforms and vendors (e.g., IBM's Office-Vision)? Do you see transparent information interchange? Please call a timeout from the quarterly earnings reports and reorganizations.

Of course, get your current house in order. Get your product line sorted out. Drop your prices across the board so that new Mac users don't keep asking me, "Why is a Mac SE so much more expensive than a 286 DOS machine, Don?" Then tell us what lies in the future. And tell us how you'll take us there. That vision made the Mac popular, successful, and different. Don't blow it now.

Third-Party Visions

If Apple isn't sure about what to tell us, it should look for inspiration among third-party vendors. It can start with its own subsidiary, Claris. The people at Claris have a real vision of the future that's not

based on Knowledge Navigator or the various ways that multimedia will enhance our lives. Claris's future is based on a careful consideration of where the Mac technology is now and what can be done to extract the most out of it in the future.

If you aren't familiar with Claris, you should be, even if you don't care one bit about the Mac. Claris has figured out that its products have to talk with each other, transparently.

Big deal, you say? Well, it certainly can be: Try to make dBASE Mac talk with FullWrite Professional. Or try to make Excel talk to Word. Sure, you can exchange parts of files, and even have those parts auto-update. But transparently? No. I'm not singling out these applications as bad; it's just that people sometimes try to use one application's data (that Excel spreadsheet) in another one (adding the spreadsheet figures to a
continued

quarterly business report in Word).

Claris has publicly (the only place it counts for its customers) committed to making its applications talk together seamlessly. Its XTND file format is part of that commitment. Its devotion to Apple's Inter-Application Communication (IAC) is another. And its product development strategy is the third.

The bottom line is that Claris is well aware that connectivity really means transparent data exchange and update among applications, Macs, and, eventually, different workstations running other operating systems. Claris envisions you sitting in front of your computer and working with the information rather than worrying about the tool. That's a real computing future that makes for a real computing strategy. I hope Apple listens.

Foresight at Farallon

Claris isn't the only third party Apple should be listening to. Farallon Computing, manufacturer of PhoneNet, is another. I recently spent a day at Farallon's new digs in Emeryville, California. The first thing that struck me was the company's knowledge about computing in

general and the Mac specifically. Farallon knows its own products (inside and out), its competitors' products, what is happening in the computer market, and where the company is going.

Although Farallon makes many dollars on its easy-to-buy-and-use Apple-Talk cabling systems, it doesn't think in terms of being a networking company. Like Claris, Farallon sees the big picture: the reason for technology, and what it should be doing for us.

As I reported in March, Farallon's new product, MediaTracks, is hot stuff. It's exactly what I mean when I say that the company has a handle on the big picture. MediaTracks is not the be-all and end-all of multimedia applications. It's not MacroMind Director or Studio/8. It allows you to edit Mac video (i.e., Mac screens you've recorded with Screen-Recorder), annotate it, add sound and voice-overs, and use it. It accomplishes this task simply and straightforwardly, the way a computer application ought to.

Farallon knows that the future of computing revolves around this kind of almost embarrassing simplicity that works across computers, computers that can be

across the country. It knows that Mac technology should be driving this realization and leading the way to real communications connectivity (which goes far beyond Apple's connectivity mantra). But don't take my word for it. Watch Farallon closely over the next year.

I just hope that Apple watches, too.

Tip of the Month

If you own a Mac Portable, be careful with the Sleep command that's executed from the Finder's Special menu or from the Battery desk accessory. It can cause some problems with files that are left open when it's executed. Sleep effectively turns off all the Portable's subsystems, except the static RAM, keeping your current application and its attached file ready for immediate use. Just press any key, and the Portable resumes where it left off, without rebooting.

All of that's fine—in theory, anyway. But I've recently discovered some problems caused by Sleep. I was working with my Mac Portable in Cupertino last week, and I had two applications (and their attached documents) open under

continued

Circle 174 on Reader Service Card

MultiFinder. I saved the documents to the hard disk and decided to get some sleep (the human kind). I was planning to get up at seven and finish my work, so I used Sleep to keep things just a keystroke away.

Awaking at seven, I pressed a key and the Portable came back to life, putting me back into MultiFinder with my two applications (VersaTerm-Pro and Nisus), just where I had left them. Except...I switched to the Nisus file. It was, to use a popular expression, toast. The screen was filled with garbage—no text. The application was OK, but the file was definitely not. I closed the file and reopened it with Nisus. Same problem. The saved version from disk was also history. Fortunately, the disk backup was OK.

After some hours of experimentation with Sleep and other Nisus test files, I duplicated the problem repeatedly. It comes down to a problem with Nisus's auto-save feature (which can be set to different timings) conflicting with Sleep. Nisus probably thinks the machine is in an active state during Sleep (which it isn't) and tries to auto-save the file, which is impossible since the hard disk is powered off.

I tested Sleep in other contexts. After doing a few experiments with MacWrite II, Excel, FileMaker II, VersaTerm-Pro, WingZ, and StuffIt, I found that Sleep can definitely trash files when these applications are left open under Multi-Finder and the command is executed. I even managed to trash the VersaTerm-Pro application itself, with an alarming ease.

I reported these problems informally to friends at Apple and found they've had problems, too, with their own Portables. My advice: Don't use Sleep to save a system context where you have multiple applications open under MultiFinder. Save your files and close the applications before invoking Sleep. Better yet, just invoke the Finder's Shut Down command, which accomplishes the same thing and turns off the Portable. For now, when I go to shut down my brain in the wee hours, I'll shut down my Portable, too. Sleep apparently is not as good. ■

Don Crabb is the director of laboratories and a senior lecturer for the computer science department at the University of Chicago. He is also a contributing editor for BYTE. He can be reached on BIX as "decrabb."

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

If you want to see a reliable computer, look for the holes in a Tatung.

Ask your Tatung dealer to show you the very quick 386-class Tatung 8000.

Turn it over. And on the bottom plate you'll notice two mysterious holes. What are they for? Well actually they're *bolt holes,* provided so you can bolt your Tatung 8000 to a combat jeep. Several armies do.

The 386SX-based Tatung 8800S. It's available now.

Of course, you may not think battlefield reliability is completely necessary at the office. Let's say your primary interest is economical performance.

Then check out our new 386SX™-based Tatung 8800S. At 2.75 MIPS, the Tatung 8800S delivers uncommon speed and power. And because it's a 386-class machine, it runs all of the latest 386-specific applications.

But at a 286-class price. What more could you ask?

For engineering at its best, see your Tatung dealer.

Call toll-free 1(800) 827-2850, ext. 253, for complete information. Tatung Company of America, Inc., 2850 El Presidio St., Long Beach, CA 90810.

Novell® and SCO® certified for file server and multi-user UNIX™ server applications. All trademarks are the property of their respective manufacturers.

TATUNG®
Engineering at its best.

Circle 268 on Reader Service Card

DIGGING INTO HPFS

Putting the new HPFS under the magnifying glass

I have now had a few months to play with the High Performance File System. As I've mentioned before, it's a bit faster than the old file-allocation-table (FAT) system, on balance, and is also chock-full of new features. Microsoft has been pretty close to the vest about the innards of HPFS, mainly because it doesn't want to give anything away that would make the patent process more difficult. Since Microsoft hasn't released some details yet, I got out my Norton Utilities and did a little spelunking. To make my examples simpler, I am assuming that all files are contained in the root directory of the disk.
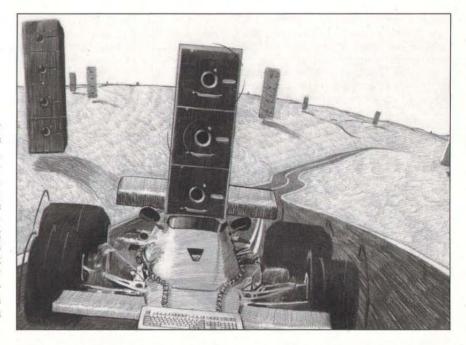
HPFS Disk Organization

Under the FAT system, the disk had the following:

- a DOS boot sector with disk ID information and a pointer to the FAT
- a FAT that indicates where your files exist, with data areas allocated in terms of multiple-sector minimum units called *clusters*
- a root directory to tell you what files you have
- a large data area for the files themselves

Under HPFS, a disk is divided up quite differently (see figure 1). It now has the following:

- A boot area with disk ID information (16 sectors).
- The Superblock (right after the boot area), which contains 32-bit pointers to major system areas, including the root directory, a map of free space on the disk, and a list of bad areas on the disk, as well as other system information (e.g., the last time that CHKDSK/F was run on the disk).
- The Spareblock (right after the Superblock), which contains other system information (I will not use any of its information here).
- Data bands. The 8 megabytes following the Spareblock is data band 1. Minimum disk space allocated to a file is no longer the cluster, but rather the smaller (less wasteful) sector. Disks are divided into these 8-MB data bands regardless of disk size.
- Free-space bit maps. Immediately following data band 1, four sectors (2K bytes) with a bit map describe the free space in data band 1. For each sector in data band 1, there is a bit. If the bit is equal to 1, the sector is occupied; if it's 0, the sector is free. Right after the free-space bit map for data band 1 is the one for data band 2. Data bands 2 and 3 sit

next to each other, allowing a single file to be as large as 16 MB in contiguous length. Thereafter, pairs of free-space bit maps are wedged between pairs of data bands. Figure 1 shows a disk layout for a 32-MB hard disk drive.

How HPFS Finds Files

Under FAT, a pointer from the directory to the FAT leads in turn to other FAT pointers, which together describe which clusters make up the file. Under HPFS, it's a mite more complex, but the net effect is faster file access. To find a fairly contiguous file in the root directory, HPFS executes the following algorithm.

First, look in the Superblock at byte offset 12 for a 4-byte value that points to something called an *fnode* (pronounced "eff node") for the root directory. That 4-byte value is a sector offset. Thus, if the value is, for example, 700, the item being looked for is in the 701st sector on

continued

the partition (701st because counting starts from 0). All "pointers" in HPFS are 4-byte sector offsets.

Next, use the Superblock pointer to find the fnode of the root directory. Fnodes are sectors that contain basic information about a file or directory. For example, a 100-sector file would actually take up 101 sectors on the disk: 100 sectors for its data and a sector for its fnode. Fnodes can take several forms, depending on how large a file is and how many extended attributes it has, but basically fnodes tell you where a file is and describe its EAs. The 72nd byte in the fnode of the root directory contains a 4-byte pointer to the root directory's sectors. In my case, the sectors ended up near the end of data band 1.

Use that 72nd byte in the root directory's fnode to access the root directory. The root directory gets allocated four sectors, as do all directories. Within the sectors, directory entries are variable-length, as the filenames can be as large as 255 characters. I'll explain what I've decoded of the directory in a minute, but it appears that each directory entry consists of 31 bytes and a filename. The di-

rectory then ends with a dummy record. Four of those 31 bytes point to an fnode for the file that you're looking for. (The file gets an fnode, just as the root directory did.)

After this, use the fnode pointer in the directory to find the fnode for the file in question. The fnode for a file describes the space that the file uses in one of two ways. The much more common way describes a file's location through the use of *sector runs*, which are 8-byte structures consisting of a 4-byte run length and a starting sector offset. For example, say the file occupies sectors 100 to 106 and 200 to 201. This file has two sector runs: a seven-sector run that starts in sector 100, and a two-sector run that starts in sector 200.

There's room for eight sector runs, so any file that's fairly contiguous (i.e., eight fragments or less) will fit nicely into this scheme. Figure 2 contains a partial fnode illustrating the eight sector runs for this file.

For the pathologically fragmented files, HPFS also allows an fnode to point to up to 12 *allocation sectors*, each of which can contain 40 sector runs of its

own. This is probably about as much as anyone could want, although HPFS can have the allocation sectors contain not sector runs but more pointers to allocation sectors, and so on.

I found that the eight length/starting sector pairs started at byte offset 40 in the fnodes that I examined.

Finally, use the sector runs to find the file itself.

HPFS Directories
What about directories? It is known that a directory entry will contain an fnode pointer, a filename, size, and several dates (i.e., date created, date last read, and date last modified). It appears that a directory entry looks like 31 bytes and a filename. Here's what I've pieced together. First, the directory sector contains 56 bytes of some kind of preamble.

continued

Figure 2: *Fnodes can point to eight sector runs directly, and many more indirectly. For all but the most pathologically fragmented files, the direct scheme should do.*
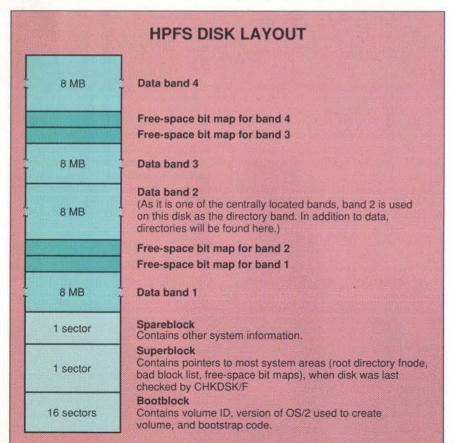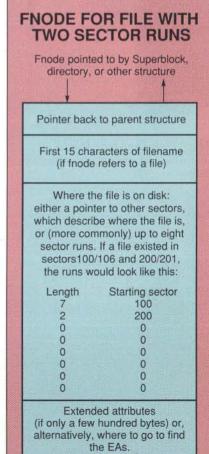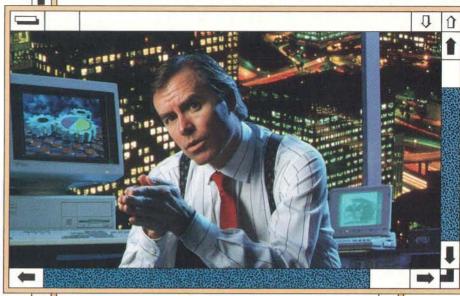
Figure 1: *Paired data bands can store 16 megabytes of contiguous sectors.*

HPFS DISK LAYOUT

8 MB	**Data band 4**
	Free-space bit map for band 4
	Free-space bit map for band 3
8 MB	**Data band 3**
8 MB	**Data band 2** (As it is one of the centrally located bands, band 2 is used on this disk as the directory band. In addition to data, directories will be found here.)
	Free-space bit map for band 2
	Free-space bit map for band 1
8 MB	**Data band 1**
1 sector	**Spareblock** Contains other system information.
1 sector	**Superblock** Contains pointers to most system areas (root directory fnode, bad block list, free-space bit maps), when disk was last checked by CHKDSK/F
16 sectors	**Bootblock** Contains volume ID, version of OS/2 used to create volume, and bootstrap code.

FNODE FOR FILE WITH TWO SECTOR RUNS

Fnode pointed to by Superblock, directory, or other structure

Pointer back to parent structure

First 15 characters of filename (if fnode refers to a file)

Where the file is on disk: either a pointer to other sectors, which describe where the file is, or (more commonly) up to eight sector runs. If a file existed in sectors 100/106 and 200/201, the runs would look like this:

Length	Starting sector
7	100
2	200
0	0
0	0
0	0
0	0
0	0
0	0

Extended attributes (if only a few hundred bytes) or, alternatively, where to go to find the EAs.

"It's great to have a 386SX system that's ready for the future.

"And it's even more appealing when I can save $200 right now."

ZENITH DATA SYSTEMS INNOVATES AGAIN™

The SX Appeal deal—$200 instant savings on our Flat Technology Monitor when you buy our Z-386 SX or SupersPort® SX.

SX Appeal™

Why does the Intel386SX™ microprocessor have so much appeal? It not only handles today's advanced applications, but assures compatibility with emerging software designed for graphical user interfaces— all with the affordability of a 286-based PC. And now for a limited time only, you can take advantage of our special *SX Appeal* deal.

Choose our Z-386 SX desktop PC that maximizes 386SX performance for ultra-fast processing speed. Or our SupersPort SX laptop PC with its *Page White* VGA display. Then for each system purchased, you'll receive **$200 savings** on our award-winning Flat Technology Monitor—the revolutionary non-glare VGA color monitor with unsurpassed clarity that's the perfect match for graphical interfaces. This offer also applies to all our 386-based desktops and VGA laptops.

To take advantage of our *SX Appeal* deal, just pick up a $200 Savings Certificate at your participating Zenith Data Systems Medallion Reseller by **June 1, 1990.** For the location nearest you, call:

1-800-227-4617
U.S. and Canada

ZENITH
data systems

Groupe Bull

© 1990 Zenith Data Systems

Graphics simulate Microsoft® *Windows*, a trademark of Microsoft Corp. Intel386SX is a trademark of Intel Corp.

Circle 300 on Reader Service Card

LAYOUT OF A DIRECTORY ENTRY

The 31 bytes preceding each filename.

Offset	Size	Description
0	4	Unknown
4	4	Pointer to file's fnode
8	4	Date of last modification
12	4	Size in bytes
16	4	Date last read
20	4	Date created
24	6	Unknown
30	1	Length of filename
31	1-255	Filename

Then come 31 bytes preceding each filename, laid out as shown in the table.

The directory is again finished with a dummy record that is all zeros except for length (1) and filename (hexadecimal FF). The new dates—date last read and date last modified—have the potential to be extremely useful.

But dates under HPFS are different than under DOS. DOS gurus know that directory dates and times are expressed so that the first 2 bytes are the time and the last 2 bytes are the date. Under HPFS, it's just a single 4-byte entry that represents seconds since 1-1-1970. Why the new format? It simplifies greatly the task of doing date arithmetic. If all dates are merely "seconds since January 1, 1970," you can calculate differences between dates by subtraction.

A Welcome Improvement

Everything in HPFS seems to have been conceived in multiple levels to provide maximum speed *and* flexibility. For example, consider the EAs. They will be useful because they let you attach up to 64K bytes of information to a file—information *about* the file. One person at Microsoft told me that he has a program that computes file checksums, which are then stored as an EA for the file. He then has a program that runs in the background computing file checksums and comparing them to the values in the EAs—OS/2 antivirus protection!

EAs sound good, but where are they stored? Microsoft seems to have reasoned thus: In most cases, there will be no more than a few hundred bytes' worth of EAs, although there may be more in a few cases. HPFS was therefore designed to store the EAs right in the fnode if they don't take much space, or (if they do need a lot of space) to move them out to a separate sector. The result: Getting to EAs in most cases is instantaneous, because HPFS is already examining a file's fnode. But for those who need it, larger EAs can be had, at a minor cost in speed.

Microsoft has designed a new system with the future in mind. The space constraints of FAT systems are gone. The indirection offered by the fnodes makes adjusting file location or size much simpler, and the fnodes contain some redundant items that FAT systems lack, which makes data recovery a bit easier. Extra dates in the directory will make backups and file maintenance simpler. All in all, HPFS is a nice improvement. ∎

ACKNOWLEDGMENT

Thanks to the Microsoft folks who let me in on the new date format and confirmed what I had figured out about the HPFS structure.

Mark J. Minasi is a managing partner at Moulton, Minasi & Company, a Columbia, Maryland, firm specializing in technical seminars. He can be reached on BIX as "mjminasi."

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

BYTE

Inside: Bonus Graphics Supp

33 MHz

SIA's new
386/32

The polar bear is the great white hunter of the Arctic Circle. Migrating long distances over the vast frozen desert, it tracks large prey such as seals and sea lions. In a sprint, it can reach speeds of up to 25 mph. When it raids the ice box, it doesn't waste energy on small fish. The readers of BYTE magazine are just as selective. For they, too, are relentless hunters with big appetites. They are advanced personal computing experts whose hunger for new product information is insatiable. They seek technical analyses. Extensive reviews. A complete rundown of product comparisons. Which is what they get in every BYTE. Some publications may promise you more paid readers than BYTE's 450,000. But none so voracious.

McGraw Hill

SHORT TAKES

ScanMan Model 32

NCR 20-MHz 386SX

dBASE IV 1.1

HP 48SX Calculator

Pivot Display

What the ScanMan Can Scan

As hand-held scanners become more popular, manufacturers are scrambling to enhance their features. Logitech, one of the many firms selling hand-held scanners, is no exception, and it has completely redesigned its ScanMan Macintosh scanner with the most common uses of these hand-held scanners in mind.

Most users of hand-held scanners are interested in easily and rapidly scanning smaller pictures or graphics of some sort and columns of text, rather than full pages, which are better handled by the more traditional flatbed and page-feeding scanners. Typically, the scanned results will be used in desktop publishing applications for newsletters and brochures, with output going to a 300-dot-per-inch laser printer. Two other common uses are scanning columns of text and numbers for use with optical-character-recognition packages and scanning images for use in software, particularly HyperCard.

The **ScanMan Model 32** scans up to 32 levels of gray, as well as black and white, allowing more accurate photographic reproduction. True gray-scale scanners store the gray information as gray levels rather than as patterns of black-and-white dots. This lets them represent the shades accurately when images are resized and edited.

I tried a beta version of the scanner. The hardware had been finished, but some bug fixes were still being made to the software. All the major functions of the program worked, however. Setting the scanner up involved plugging it into the chain of SCSI devices using its control box and then installing some system software. The system software can be run as a desk accessory or as a stand-alone program. The scanner has a new case, which lets scans be much straighter and lets you see more of what is being scanned.

Unlike most hand-held scanners, which rely on a roller near the scanner head to help the scanner move across the paper and measure the speed at which it is being moved, the ScanMan uses a combination of the roller and

two extra wheels on the lower part of the device. These help keep you from twisting the scanner sideways while scanning, resulting in a straighter and more accurate scan. This proved to be very smooth to roll and had a reasonable resistance to being moved sideways. Another nice innovation is a scanning speed LED on the case. This lit up to let me know when I was scanning too fast and losing image data.

The ScanMan now has a new scanning head with greater accuracy. It is a true 400-dpi scanner and includes software to let it scan up to 32 levels of gray. This is achieved by having the scanner use an area six dots square to measure for gray levels. This technique reduces the accuracy of the scanner to 66 dpi when using 32 gray levels. However, laser printers can also print only gray scales by using shading patterns of black-

and-white dots. The result is that a gray-scale scan made using the Logitech scanner will print out perfectly on a laser printer. However, it will not look any better if output goes to a typesetting machine like a Linotron. Certainly it was possible to get very good output from scanned images printed on a laser printer and even on an ImageWriter, although care was needed here in the exact choice of resolution and halftoning method.

The scanning software provided with ScanMan now lets you scan horizontally or vertically by simply clicking on an icon. Resolution can be varied from 100 dpi up to 400 dpi by setting a switch. Brightness and line art or gray-scale select switches are also included. These settings are recognized automatically by the scanner, so you don't have to choose them from software as well.

The software has a basic but full set of image-editing tools. Most important is the ability to save copies of images in a wide range of formats.

File formats supported include 1- or 8-bit TIFF, PICT or PICT2, and MacPaint. You can also select different resolutions depending on whether output is to go to the screen, an ImageWriter, or a laser printer. The software lets you perform halftoning in several ways, since different techniques work better for different purposes.

The ScanMan is a carefully thought-out extension to previous hand-held scanners. Its ease of use and functionality have both been considerably enhanced, and I highly recommend it when typeset quality is not required.

—Owen Linderholm

NCR's 386SX: Quick, Quiet, and Secure

As good an example as any of the quality of German engineering, NCR's new small-footprint **20-MHz 386SX** offers speed, comfort, and security.

The motherboard includes 2 megabytes of RAM (expandable to 32 MB), an on-board Super VGA controller, two serial ports, and a parallel printer port. Some "special" features are a 16K-byte memory cache using "two-way anticipated caching" methods and an ISA bus-write buffer. The back of the motherboard has an additional 37-pin connector for an external hard disk drive or streaming tape drive. The keyboard jack is the small IBM PS/2 variety. The case and chassis wrap all the components in steel and lock them with a solid tubular key lock for hardware security, low noise, and RF shielding. The system is powered by a 175-watt auto-sensing power supply that doesn't require any configuration for voltage or frequency. A nearly silent fan cools the power supply, as well as the rest of the system.

The hard and floppy disk drives are mounted on a separate internal subsystem that is easy to remove and install. They do not require additional controller boards since they have their own controller hardware. This means that all five ISA bus slots are available for peripheral boards, even when the system is configured with three drives.

NCR shares the lead in the commercial and banking transaction market, and it has

very few competitors. This machine is a good example of why: Although it is a general-purpose SX, it has features that make it especially attractive for the large office. It is as quiet as a laptop. But more important, it has security as

its central theme. A secure system with secure software makes it difficult for unauthorized users to access the system. This computer has an embedded security system with a master password and three user-level passwords. If

the security feature is used, the system will not boot without a password. Only the master password can reconfigure the system with the setup menu. Applications software can access calls to the security system to evaluate what level of security the user possesses.

You might be able to get around all this by installing your own routines by means of the floppy disk drive. But a secure system would not have a floppy disk drive. This configuration is the reason for the floppy disk drive connector at the back of the machine. You can plug an external drive into this to do system maintenance, backups (onto the streaming tape), or system builds. This port, as well as any of the communications ports, can be disabled from the system setup menu.

We evaluated the machine with both MS-DOS and SCO Xenix/386. The MS-DOS CPU benchmarks showed that this machine runs nearly 75 percent faster than most 16-MHz 386SXs and as fast as most 20-MHz 386 machines. Even using the 32-bit benchmarks of Xenix, the new NCR 20-MHz 386SX showed between 70 percent and 90 percent of the performance of a full 32-bit system with a 20-MHz 386 processor and zero-wait-state memory. Very impressive.
—Ben Smith and Roger Adams

A Tuned-Up dBASE IV

Ashton-Tate has tuned up dBASE IV and made it a little easier to use in its new incarnation as **dBASE IV 1.1.** It now requires only 450K bytes of RAM (instead of 512K bytes), a disk-cach-

ing program is included in the new version, and some performance-tuning options are available to speed up operations.

Version 1.1 has its own disk-cache code for single-

user operation. The cache is installed automatically into all available extended or expanded memory when you load the program. A command-line option lets you
continued

limit the amount of memory used, but this is hardly needed since the cache is unloaded automatically when

the program is closed, freeing the memory. The cache does speed things up. For example, sorting a database of 1419 10-field, 173-byte records on a single field took 14 seconds not using the cache and 7 seconds using it. Sorting a larger database, 3230 records with the same specifications on a single field, took 29 seconds without the cache and 19 seconds with it.

The performance-tuning options are BAK, TMP, and DBHEAP. Since they're DOS environment variables, they must be set from DOS before loading dBASE. BAK and TMP specify paths to backup and temporary files. The DBHEAP allows the juggling of RAM allocation between dot prompts and overlay swapping. You can optimize

the execution of dBASE programs by tweaking this memory allocation to fit the program you're using.

Ashton-Tate has removed many of the version 1.0 restrictions on recursive user-defined functions (UDFs) and all the ON commands. A new command, DBTRAP, was added to protect you from shooting yourself in the foot. When DBTRAP is on, the system takes steps such as saving parts of the existing environment, prohibiting certain recursions, and blocking UDFs from using certain commands (e.g., CLOSE, PACK, and MODIFY STRUCTURE).

Be aware that dBASE IV isn't something you install and immediately begin using to write complex reports drawing on multiple data

tables. But if you've used earlier incarnations of dBASE, you should feel right at home, although you'll have to read the manual to take advantage of many of the functions and features. The familiar dot-command structure is there for dBASE programmers. For novices and casual users, there is the Command Center, a point-and-shoot facility for performing most functions. If you're totally new to dBASE, for a program as large and complex as this, it's quite easy to use. You'll face somewhat of a learning curve, but it's gentle and easily managed. And the choice of using query by example or Structured Query Language to generate reports is a real plus.
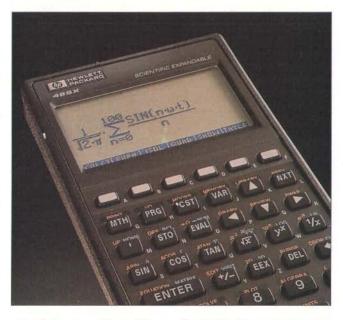
—George Bond

HP's New Calculator Is Virtually a Hand-Held PC

Hewlett-Packard's HP 41 calculator has just reached the end of its 10-year life cycle and is being replaced by the **HP 48SX**. The original HP 41 was basically a calculator for scientists and engineers, and it included serial I/O, a full programming language, and memory expandability.

The 48SX is more than just a calculator. The base model comes with 32K bytes of RAM and 256K bytes of ROM. It has an RS-232C serial port and two slots for RAM or ROM cards, letting you expand memory up to 288K bytes or use add-on programs like equation libraries for specific applications. An optional HP Solve Equation Library card is available, which includes over 300 equations for engineering and scientific problems ranging from columns and beams to fluid flow or magnetism.

HP is offering an optional serial interface kit with software, which links the

calculator to either Macintosh or DOS systems. In fact, the kit lets you control the calculator from the host computer, allowing you to use the calculator as an equation-solving peripheral. The calculator also has a two-way infrared interface for use with HP's infrared printer or

for exchanging programs or data with another 48SX.

The 48SX fits in the palm of your hand and weighs only 9 ounces. Unlike the 28S, the 48SX does not have the folding-notebook-style design, which was unpopular with HP customers.

In addition, the 8-row by

22-column display is greatly improved over the display of the 28S. The new calculator uses a supertwist LCD design, allowing better contrast and readability. The HP 48SX has impressive graphics and calculus functions, finding roots, intersec-
continued

*Embedded systems designers have already used **CrossCode C** in over 577 different applications.*

How to choose a 68000 C compiler for your ROMable code development

These twelve important **CrossCode C** features could make the difference between success and failure

It's hard to know ahead of time what features you'll be needing in a 68000 C compiler. But if you're using **CrossCode C** you won't need to think ahead, because **CrossCode C** is already equipped with these twelve important features for your ROMable code development:

1. A 100% ROMable Compiler: **CrossCode C** splits its output into five memory sections for easy placement into ROM or RAM at link time.

2. Integrated C and Assembler: You can write your code in any combination of C and assembly language.

3. Readable Assembly Language Output: The compiler generates assembly language code *with your C language source code embedded as comments*, so you can see each statement's compiled output.

4. Optimized Code: CrossCode C uses minimum required precision when evaluating expressions. It also "folds" constants at compilation time, converts multiplications to shifts when possible, and eliminates superfluous branches.

5. Custom Optimization: You can optimize compiler output for your application because *you* control the sizes of C types, including pointers, *floats*, and all integral types.

6. Register Optimization: Ten registers are reserved for your register variables, and there's an option to automatically declare all stack variables as *register*, so you can instantly optimize programs that were written without registers in mind.

7. C Library Source: An extensive C library containing over 70 C functions is provided in source form.

8. No Limitations: No matter how large your program is, **CrossCode C** will compile it. There are no limits on the number of symbols in your program, the size of your input file, or the size of a C function.

9. 68030 Support: If you're using the 68030, **CrossCode C** will use its extra instructions and addressing modes.

10. Floating Point Support: If you're using the 68881, the compiler performs floating point operations through the coprocessor, and floating point register variables are stored in 68881 registers.

11. Position Independence: Both position independent code and data can be generated if needed.

12. ANSI Standards: CrossCode C tracks the ANSI C standard, so *your* code will always be standard, too.

There's More

CrossCode C comes with an assembler, a linker, and a tool to help you prepare your object code for transmission to PROM programmers and emulators. And there's another special tool that gives you symbolic debugging support by helping you to prepare symbol tables for virtually all types of emulators.

CrossCode C is available under MS-DOS for just $1995, and it runs on all IBM PCs and compatibles (640K memory and hard disk are required). Also available under UNIX, XENIX, and VMS.

CALL TODAY for more information:

1-800-448-7733
(ask for extension **2001**)

Outside the United States, please dial

PHONE: **1-708-971-8170**
FAX: **1-708-971-8513**

SOFTWARE DEVELOPMENT SYSTEMS, INC.
DEPARTMENT 21
4248 BELLE AIRE LANE
DOWNERS GROVE, ILLINOIS 60515 USA

CrossCode™ is a trademark of SOFTWARE DEVELOPMENT SYSTEMS, INC. MS-DOS® is a registered trademark of Microsoft. UNIX® is a registered trademark of AT&T. XENIX® is a registered trademark of Microsoft.

tions, derivatives, slopes, and so on.

The 48SX has several new features that engineers and scientists will find very useful. Unlike previous scientific calculators from Hewlett-Packard, the 48SX lets you actually enter an equation the way you would with paper and pencil. Previously, you could only store

variable values in the calculator and then execute the equation by translating it into keystrokes. With the 48SX, you can type in the actual equation, supply the variable values, and hit the Solve key.

Another feature that I found impressive is the automatic-unit management capability of the 48SX. The calculator can automatically

convert units into the resultant units that you specify. For example, you can enter meters and seconds and tell the 48SX to solve the equation in miles per hour. Or you could enter newtons, meters, and hours, and have the calculator solve the equation in watts.

The 48SX calculator also comes with a new high-level

programming language. It is not compatible with the language used in the HP 41, but an emulator is provided so that HP 41 programs will run on the 48SX.

Overall, I found the Hewlett-Packard 48SX to be an impressive calculator for serious engineering and science applications.

—*Nick Baran*

A Mac Monitor That's Sure to Turn a Few Heads

P **ivot** is a Macintosh full-page monitor that offers both portrait (vertical) and landscape (horizontal) orientations. What's unique about it is that you can twist the

Pivot display a full 90 degrees, going from full-page portrait to a full-page landscape mode, and the image on your screen will rotate as well.

Pivot
$1690

Requirements:
Mac SE/30 or Mac II family.

Radius, Inc.
1710 Fortune Dr.
San Jose, CA 95131
(408) 434-1010
Inquiry 999.

If you need a full-page display but also need both width and length for different applications, Pivot might be just the ticket.

Pivot works with all existing Mac software, Radius claims. Its 15-inch diagonal screen size offers a portrait resolution of 640 by 864 pixels and a landscape resolution of 864 by 640 pixels. I found the display clear, crisp, and impressive, as I would expect from a full-page monitor. The Pivot weighs slightly more than Radius's standard full-page display, but its dimensions are almost identical.

Pivot displays two or four shades of gray and is upgradable to 16 shades. It works with either a Mac SE/30 or any member of the Mac II family. Pivot works as a standard full-page display (portrait only) with either a Mac Plus or SE.

Designing Pivot's unique rotating display created many headaches for Radius's engineering staff. First, they had to develop a high-speed pixel rotation engine for the interface card. This they did, creating one that rotates pixels 90 degrees at a rate of over 51 MHz.

Indeed, I found Pivot very fast in redrawing a Microsoft Excel spreadsheet. For example, when I tilted the display from portrait to landscape, the screen would go blank for about a second and then the spreadsheet would reappear, including the cells

that had been left out of the portrait display. The speed of the display redraw was the same—regardless of whether I was going from portrait to landscape, or vice versa. Also, I found the display easy to rotate with just a finger.
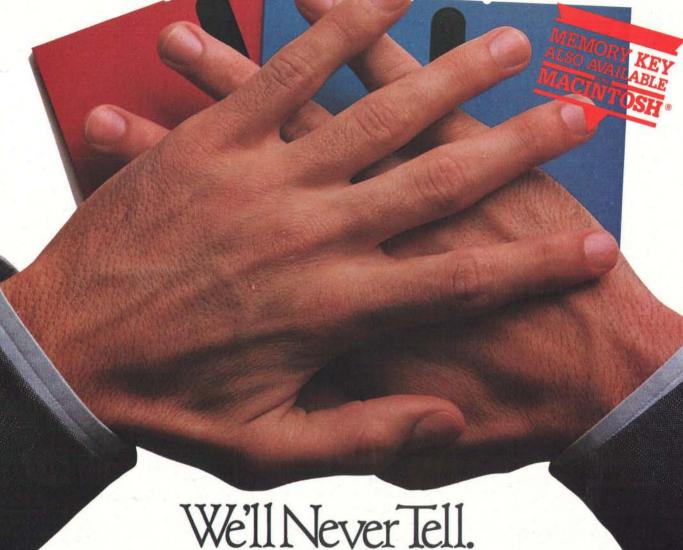
Pivot includes a "position-sensing device" that sends an interrupt to the computer whenever you change the display's position. The monitor's system software then reprograms the bit-map orientation, reorganizing the Mac Desktop to match either portrait or landscape mode. Also interesting is that Pivot scans in the same direction in both modes: from left to right, top to bottom in portrait; and from left to right, bottom to top in landscape.

Pivot's internal circuitry is positioned to allow proper cooling in either landscape or portrait mode. Similarly, the monitor's magnetic shield is designed to work in either.

The Pivot's list price is $1690 (which includes the Mac SE/30 or NuBus interface board), making it actually only $200 more than Radius's standard full-page display monitor. Sure, full-page displays aren't cheap, but it's nice to see Radius offering the unique Pivot at a competitive price. And while not every full-page monitor user needs the extra functionality of landscape mode, those who do might swivel their heads to take a closer look at Pivot. ∎

—*Jeffrey Bertolucci*

Commodore Sets Course for Multimedia

■ Bob Ryan

I magine the ideal multimedia platform. It would have impressive graphics and sound, be easy to use, and work transparently with sophisticated video equipment. Throw in a low-overhead multitasking operating system, and you'd have all the pieces you need to run sophisticated multimedia applications. You'd also have an Amiga from Commodore.

Realizing this, Commodore has finally defined a focus for the Amiga line and staked its claim to the emerging multimedia market. The centerpiece of the company's strategy is a new, 68030-based computer, the Amiga 3000.

The Amiga 3000 (see photo 1) has many impressive features, the best of which may be its price. With the 16-MHz model selling for $3495 and the 25-MHz version fetching a mere $3995, the new system costs nearly $3000 less than a comparably equipped Mac IIci. This aggressive pricing, combined with the new look and capabilities of the machine, makes the Amiga 3000 a formidable entry into the multimedia marketplace.

Bits and Pieces
The Amiga 3000's two models are the 3000/16, a 68030/68881-based machine

The Amiga 3000 boasts

enhanced custom chips,

a fast 32-bit processor,

and a stunning new look to

its Workbench interface

running at 16 MHz, and the 3000/25, a 68030/68882 machine that runs at 25 MHz. Except for clocks and coprocessors, the models are identical.

The Amiga 3000 is the first Amiga with a native 32-bit architecture. With the exception of the graphics system, which uses a 16-bit bus and a National Television System Committee (NTSC)–compatible 7.14-MHz clock, all the internal pathways in the Amiga 3000 are 32 bits wide and use the processor's clock. Commodore has also managed to make the external I/O bus capable of 32-bit transfers without losing compatibility with 16-bit peripheral cards. The I/O bus—called Zorro 3—multiplexes the address and data lines for I/O cards capable of 32-bit transfers, while treating 16-bit Zorro 2 cards normally.

Goodbye, Flicker
The heart of the Amiga 3000 is an enhanced version of the custom chip set that defines the Amiga line. The enhanced Denise chip can output up to 1280 horizontal pixels in four colors, and it sports a new 640- by 480-pixel-resolution productivity mode. (Unlike other high-resolution modes on the Amiga, the productivity mode is not interlaced, and it is limited to four colors.) Being programmable, Denise lets you mix different horizontal and vertical resolutions and provides the overscanning necessary for video applications. The new Agnus chip doubles the amount of accessible chip RAM to 2 megabytes. (In Amiga parlance, *chip RAM* is memory accessible by both the CPU and the custom chips. It is used primarily to hold graphics and sound data that gets massaged by the custom chips. Memory available to the CPU exclusively—and used primarily to hold executable code and nondisplay data—is called *fast RAM*.) The only chip not upgraded is Paula, which provides the Amiga's four-voice sound.

New to the graphics system is a de-interlacing/scan doubling capability,

which takes interlaced high-resolution screens and outputs them without flicker to a 31.5-kHz VGA monitor. It also eliminates the black scan lines between rows of pixels on both low- and high-resolution displays. As a result, the Amiga 3000 produces a rock-steady high-resolution display. Unlike earlier Amigas, which by default ran the Workbench interface in medium resolution (640 by 200 pixels), the Amiga 3000 can effectively use high resolution (640 by 400 pixels) without additional hardware.

Although the new capabilities of the graphics system are significant, they don't address the larger question of how to upgrade Amiga graphics to 8- and 24-bit color while maintaining compatibility with current graphics software. (The tight coupling of the graphics system and the Amiga CPU, which gives the machine its superior graphics performance, ironically also makes it more difficult to upgrade the graphics capabilities.) Palette limitations on the Amiga 3000 are the same as for earlier Amigas: 32 colors in low resolution and 16 in high resolution, of a possible 4096. Techniques such as extra half-bright, dynamic high-resolution, and hold-and-modify let you increase the number of colors on screen (up to 4096 in the case of HAM), but they either force you to swap palettes on the fly—a serious drain on system bandwidth—or don't let you define each pixel independently. (The latter drawback is actually an advantage when you use HAM mode for video work; because it can take up to three pixels to change from one color to another, HAM mode is naturally antialiasing.)

Despite these alternative graphics modes, the lack of native 8- and 24-bit color on the Amiga is a concern. Such capabilities will be necessary for next-generation video and multimedia applications and are already available for many of the Macintosh, MS-DOS, and Unix machines. The Amiga must provide for

Photo 1: The Amiga 3000 features 2 MB of memory, an 880K-byte floppy disk drive, and a 50-MB SCSI hard disk drive.

8- and 24-bit color to remain competitive. Amiga Product Manager Keith Masavage wouldn't comment on Commodore's plans in this area other than to say that "both short-term and long-term solutions are in the works."

Motherboard Tour

The Amiga 3000 motherboard reflects a high degree of integration (see photo 2). In addition to the enhanced Agnus and Denise chips, the motherboard contains eight other proprietary chips fabricated by Commodore. These include a bus controller, a RAM controller, and a DMA controller. The DMA controller works in conjunction with a SCSI controller chip to provide 32-bit DMA for

SCSI I/O devices such as hard disk drives. The Amiga 3000 uses 512K bytes of ROM to hold the multitasking executive, system libraries, and disk operating system. Commodore also plans to upgrade the Amiga 500's and the Amiga 2000's ROMs to 512K bytes.

The RAM controller—dubbed Ramsey—manages access to fast RAM on the motherboard. The Amiga 3000 comes with 2 MB of RAM—1 MB of chip RAM and 1 MB of fast RAM. It also comes with sockets that can hold up to 4 MB of fast RAM using 4 by 256K-bit paged or static-column zig-zag in-line packages, or 16 MB using 4 by 1-megabit ZIPs.

The Amiga 3000 doesn't use a large external cache to speed memory access

by the 68030. It does, however, handle both burst-mode and page-mode access to fast RAM. A standard memory access to fast RAM involves two wait states, while burst mode can result effectively in zero-wait-state performance. Under ideal circumstances, the page mode can reach transfer rates of over 33 MB per second. The Ramsey chip determines which mode is appropriate for any particular memory access.

The megabyte of chip RAM in the system comes in the familiar DIP packaging, and the motherboard contains sockets for another megabyte. The megabyte of fast RAM on the motherboard also comes in socketed DIP chips. When you upgrade your fast RAM using ZIP packs, you move the original fast RAM to the empty chip RAM sockets to give you the maximum 2 MB of chip RAM.

The Amiga 3000 comes with a standard Amiga 3½-inch 880K-byte floppy disk drive and an internal 3½-inch 50-MB 17-ms SCSI hard disk drive. It also lets you mount internally another 3½-inch disk drive—floppy or hard.

Expanding Your System

The expansion backplane on the Amiga 3000 attaches to two sockets on the motherboard. The backplane lets expansion hardware automatically configure itself during start-up and supplies DMA access to the motherboard. The backplane supplies four Zorro 3 expansion slots. In addition, the backplane has two AT-compatible slots in line with two of the Zorro slots, and a video slot in line with a third. You use the AT slots in conjunction with the BridgeBoard cards from Commodore, which provide the Amiga with XT and AT compatibility. The video slot gives expansion cards direct access to the Amiga video-output signals. It is useful for devices such as video-effects generators and genlocks.

The Amiga 3000 also supplies several standard I/O ports. The keyboard and two mouse ports are located on the right side of the machine. On the back are a parallel port, a serial port, a floppy disk drive port, a SCSI port, two audio ports, a 15.75-kHz RGB video port, and a 31.5-kHz VGA video port. The RGB port is identical to those on earlier Amigas, ensuring compatibility with earlier monitors and other video equipment. Its signal doesn't run through the deinterlace/scan doubling hardware. The VGA port, which *is* connected to the new display hardware, drives any standard VGA monitor. You can disable this hardware—or adjust its output—using two controls

Photo 2: The Amiga 3000/25 motherboard integrates many functions into discrete chips. The expansion backplane connects to the two 100-pin sockets at left-center. Expansion boards thus lay parallel to the motherboard. Note the 200-pin CPU slot at right-front. (Food for thought: The 68040 comes in a 179-pin package.)

Photo 3: *The Amiga OS 2.0 sports an all-new Workbench interface. Built on new versions of the Amiga shared libraries, Workbench 2.0 provides easy access to multiple programs. A hardware blitter enhances its operation by providing the system with a rapid way to move rectangular blocks of memory, such as those that define a window.*

on the back of the machine.

The most significant expansion connector in the Amiga 3000 is a 200-pin CPU slot. Although it's incompatible with the 86-pin CPU slot on the Amiga 2000, it serves the same function. It permits you to upgrade the 3000 with a next-generation processor—in this case, the 68040 and its descendants. You could also use it to hold a large external memory cache for the on-board 68030.

Cleaning Up Workbench

In the past, many people dismissed the Amiga as an unprofessional machine because of the look of its Workbench interface. When all you saw was a four-color, 640- by 200-pixel screen with visible scan lines between each row of pixels, it was easy to forget the sophisticated hardware and software that lay beneath it.

The Amiga 3000 uses version 2.0 of the Amiga operating system. With this version, Commodore has reworked both the look and the functionality of the Workbench interface by upgrading or rewriting most of the underlying system software. The result is a new interface (see photo 3) that can hold its own against any GUI on the market. It provides consistent and easy access to the multitasking ability of the Amiga.

Among the improvements in Workbench are enhanced Preferences tools for setting your system configuration, new Workbench menus, alternate window views, Adobe fonts, bit-map font scaling,

horizontal and vertical scrolling of the Workbench screen, a new front-to-back gadget, a new window-selection procedure, the ability to view files that lack .info files, the ability to execute Amiga-DOS commands without using the AmigaDOS shell, the ability to start an application by "dropping" a data file into it, and many new and improved utilities. In addition, Commodore has integrated ARexx into the operating system. ARexx is a macro language that lets you share data between two or more Amiga programs equipped with ARexx message ports. (Such ports have lately become de rigueur on Amiga applications.) ARexx allows you to create "meta-applications" using stand-alone Amiga programs.

Amiga OS 2.0 also has a lot of new goodies for programmers. First, Commodore is at last including standard file and font requesters with the Amiga. In addition, it is supplying a gadget toolkit. This lets you integrate custom gadgets into the Intuition windowing system. It also provides an object-oriented interface to gadget programming. Finally, 2.0 provides a "hotlink" service that notifies your program about events of interest to it. This is particularly useful when you change a Preference setting that your program must be aware of.

Commodore designed 2.0 to be compatible with software written for earlier versions of the operating system. According to Commodore, programs that abide by the rules will run under 2.0. However, Commodore knows that everyone didn't follow all the rules, and the company is working with developers to fix software that may break under 2.0.

Currently, the Amiga has a good-size software base, although it is skewed toward graphics and video applications, as opposed to mainstream business programs. Because the Amiga 500 and 2000 can be upgraded with the enhanced chip

set and run OS 2.0, programs written for the 3000 will also run on most of the Amiga installed base of nearly 1.5 million machines.

A New Beginning?

The Amiga 3000 is an impressive machine, and Commodore is supplying it with strong support. In addition to the 3000, Commodore is preparing well over a dozen new hardware and software products designed to position the machine in a few key markets. The first of these is business, where Commodore hopes the Amiga's multimedia capabilities will provide the kind of entrée that desktop publishing gave to the Macintosh. Central to the company's business push is AmigaVision, an icon-based authoring system that lets anyone create interactive multimedia applications.

Also, Commodore is finally acknowledging the importance of networking by providing software and hardware solutions, either directly or through third parties. For the short term, this means Ethernet and third-party X Window System server software (X/11 from Gfx-Base). By in the middle of this year, it will mean TCP/IP-NFS and NetWare client software. Beyond that, Commodore has plans to support ARCnet and Flash-Talk hardware and assorted network operating systems. The company also hopes that, by increasing its penetration into businesses through networking and multimedia, it will attract more mainstream developers to the system.

Commodore is also actively pursuing the higher-education and government markets, in large part by providing Unix in addition to the new networking products. It is actively porting Unix System V.4 to the Amiga. Even with the additional cost of the operating system, the Amiga 3000 will make an inexpensive, yet powerful, Unix workstation. Finally, Commodore hopes to maintain a large share of the home/personal productivity market, primarily through sales of the Amiga 500 and low-end versions of the Amiga 2000.
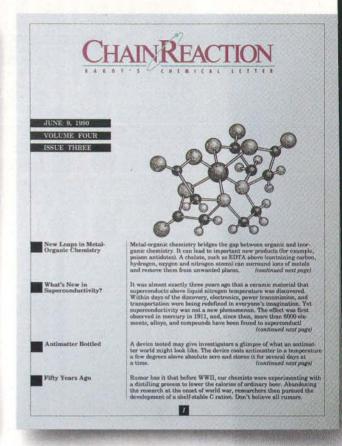
With a primary focus in multimedia and secondary pushes in education, government, and productivity, Commodore has at last outlined a strategy that takes advantage of the Amiga's strengths. With the Amiga 3000, it has produced the most capable multimedia platform you can get in a single box. If the machine were the only criterion, I'd already call Commodore's strategy a success. ∎

COMPANY INFORMATION

Commodore Business Machines, Inc.
1600 Wilson Dr.
West Chester, PA 19380
(215) 431-9100
Inquiry 1080.

Bob Ryan is a BYTE technical editor. You can reach him on BIX as "b.ryan."

Other laser printers play with one standard dot size.

HP makes it a
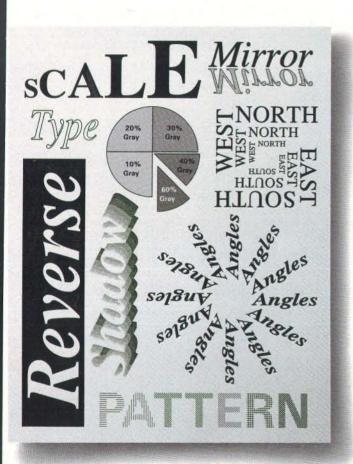
Introducing the new HP LaserJet III printer with Resolution Enhancement technology.
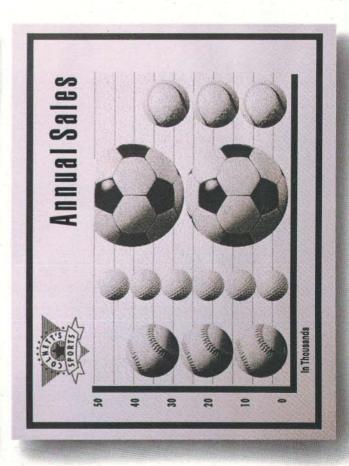
The rules have changed. Now the name of the game is Resolution Enhancement technology. You'll call it the best thing to happen to laser printing since the very first HP LaserJet printer. It gives you clearer resolution. Curves that really curve. And edges that are never jagged.

Instead of a "one-size-fits-all" dot, HP's built-in intelligence varies dot sizes. So they can fill areas where

they could never go before. For clearer, more professional-looking documents.

Resolution Enhancement technology shrinks dots to fit in curves and diagonals where they've never gone before.

whole new ball game.

But there's more than better print quality. 14 bit-mapped fonts and 8 internal scalable typefaces provide thousands of options. And enhancements to our PCL5 printer language, including our HP-GL/2 graphics language, let you print portrait and landscape on the same page. Reverse and angled type. Spirals. Even shaded text. You can also plug in Adobe PostScript® software.

For all its new features, the $2,395* list price of the HP LaserJet III is a good deal less than the HP LaserJet

Series II printer it replaces. With the same hardware compatibility, wide range of applications, 8 ppm print speed, and software compatibility, including WordPerfect 5.1 and WordStar® 6.0. And the same

reliability as the rest of the HP printer family.

So call **1-800-752-0900, Ext. 1007.** Ask for our booklet on Resolution Enhancement technology and where to find your nearest authorized HP dealer. We'll put you in a whole new league.

There is a better way.

HEWLETT PACKARD

Two new 33-MHz 486 PCs

set speed records, but

that's not the whole story

The Fast Keep Getting Faster

■ Michael Nadeau

I f speed was all ALR's and AST Research's 33-MHz i486-based PCs had to offer, I wouldn't be writing this article. They are fast—real fast. But both machines represent changes in marketing and design that could affect all Intel-based PCs.

You can buy the ALR (Advanced Logic Research) PowerVEISA 486/33, nicely configured, for a price comparable to or less than that of some 386 systems—you even get the fast Extended Industry Standard Architecture (EISA) bus. AST hasn't set a price for its Premium 486/33 yet, but its estimated list price is under $10,000. If that's too high, you can buy either box in 386 form and, with a quick board swap, later upgrade it to a full-blown 33-MHz 486 system. Clearly, the pricing of these and other 486 systems will force 386 system prices downward, and the systems' designs indicate that the ability to upgrade an existing PC is becoming an important factor to many customers.

First, the Hardware
From the outside, both the ALR PowerVEISA 486/33 and the AST Premium 486/33 look just like other high-end offerings from those companies (see photo 1). The PowerVEISA is a tower design that uses the EISA 32-bit bus. The Premium comes in desktop form with Industry Standard Architecture 16-bit slots. EISA tower and desktop configurations are available, but AST would not give a price for either. To upgrade the Premium from ISA to EISA requires a motherboard swap ($1250 for the 486/25 version). The ALR and AST systems should be available by the time you read this.

The PowerVEISA (for versatile expandable integrated scalable architecture) offers 33-MHz 386 and 25- and 33-MHz i486 CPU options. Prices range from $5795 to $8495 for the base Model 80 configuration (with 5 megabytes of RAM; 64K bytes of RAM cache; a 1.2-
continued

Photo 1: *The AST Premium 486/33 (left) is a desktop design with an ISA bus, while the ALR PowerVEISA 486/33 sports a tower design with an EISA bus. AST offers EISA versions in both desktop and tower form.*

MB floppy disk drive; an 80-MB hard disk drive and integrated device electronics [IDE] controller; five drive bays; serial, parallel, and mouse ports; and a 200-watt power supply). The 25- and 33-MHz i486 upgrade modules cost $1995 and $3195, respectively; the company will buy back the 25-MHz i486 module for $1200 when you upgrade to the 33-MHz version.

The CPU module is about half the size of AST's (see photo 2). It is home to the CPU, the FPU socket, and associated components. The cache memory card uses ALR's FlexCache architecture, which consists of 64K bytes of 25-nanosecond static RAM (SRAM).

Nine expansion slots are available, including one 8-bit and two 16-bit ISA slots, three 32-bit EISA slots, and three proprietary slots that hold the CPU, expansion memory, and RAM cache cards.

The motherboard holds up to 17 MB of 80-ns 1-MB or 4-MB single-in-line-memory-module RAM (the 32-bit memory card can hold another 32 MB of RAM). It is also home to the integrated serial, parallel, mouse, and disk I/O functions. The unit we received from ALR was a standard Model 80 with an ALR VGA monitor and graphics card.

Next!

The Premium has five 16-bit ISA slots, one 8-bit ISA slot, and three proprietary

You won't notice a difference in performance unless your applications are the most power-hungry.

CUPID (Completely Universal Processor and I/O Design) slots, which also make use of three of the 16-bit slots. The CUPID slots hold the CPU card (called the Fastboard) and 32-bit memory (16 MB for each slot). CPU options range from a 16-MHz 386SX to the 33-MHz i486. AST expects the standard configuration for the Premium to be at least 4 MB of RAM, a 5¼-inch 1.2-MB floppy disk drive, MS-DOS 3.3, an embedded hard disk drive controller, one parallel and two serial ports, and five drive bays. The system we received from AST contained 8 MB of RAM, an 80-MB hard disk drive and a Western Digital hard disk drive controller, and an AST

VGA monitor and graphics card.

AST used a different upgrade approach from ALR; the Fastboard holds all hardware dependent on the speed of the CPU, up to 4 MB of 80-ns SRAM, and a slot for an Intel or Weitek 4167 math coprocessor. The system board, identical for all CUPID models, contains the AST BIOS, expansion slots, application-specific ICs, and I/O components.

Prices for the upgrade boards were not available at press time, but they will depend on what machine you are upgrading. For instance, you will pay more for the Fastboard 486/33 for a 386SX system than for a 25-MHz 486 system, even though the boards are the same. AST's literature says that you must return your old CPU board, for no rebate, when upgrading, but an AST spokesperson said that trading your old board was up to the customer.

Both systems looked like the prototypes they are; wire jumpers populated both the upgrade modules and system boards. Both ran flawlessly, though, which is more than I can say for the first 25-MHz 486 systems BYTE encountered. Intel's experience with that chip has obviously given it a running start with the 33-MHz version.

Performance Plus

On our CPU and FPU benchmark tests, both machines blew away everything else

Photo 2: *ALR's CPU upgrade module is much smaller than the AST Fastboard card, and it uses only the separate proprietary slot on the ALR motherboard. AST's card uses one of the standard 16-bit slots, as well as a proprietary 32-bit slot. Both designs make for simple upgrades.*

BYTE BENCHMARK INDEXES

ALR's FlexCache architecture seems to be the difference in the CPU index scores. Nonetheless, both machines are screamers, especially on the FPU side. Both machines are prototypes, and these numbers should be considered preliminary.

	CPU	FPU	Disk I/O	Video
ALR PowerVEISA 486/33	9.69	37.03	3.48	4.02
AST Premium 486/33	8.21	37.10	*	3.40
Cheetah Gold 33 (25-MHz i486)	6.52	21.49	9.49	5.57
Compaq Deskpro 386/33	6.09	15.50	2.90	4.53

*Did not complete one test in our disk I/O benchmarks, for reasons we could not determine.

Indexes show relative performance; for all indexes, an 8-MHz IBM PC AT = 1. For a full description of all the benchmarks, see "Introducing the New BYTE Benchmarks," June 1988 BYTE.

BYTE has seen. The ALR PowerVEISA 486/33 easily bested the AST Premium 486/33 on the CPU index: 9.69 versus 8.21 (see the table). I attribute the wide gap in the CPU ratings to ALR's external 64K-byte read- and write-back cache. The most significant differences were in the string move tests, which benefit from a good cache. The Premium 486/33 depends on the 8K-byte cache built into the i486.

Both machines blasted previous FPU marks by about 10 index points: 37.03 for the PowerVEISA 486/33 and 37.10 for the Premium 486/33. This should make a strong argument for 33-MHz i486s in math-intensive applications, such as CAD or serious financial work.

The PowerVEISA 486/33 appears to have the edge in price and performance. Both machines are prototypes, however, and prices and benchmark numbers can change. All prices given here are retail list; dealers will usually offer significant discounts.

For the Power-Hungry Only?

In truth, you won't notice a difference in performance between these machines, or even between these machines and fast 386 systems, unless your applications are among the most power-hungry. But even everyday applications are demanding more and more processing power.

Graphical user interfaces require a great deal of computing overhead and encourage software developers to add sophisticated features. A word processor today is passé unless it can generate graphics, perform desktop publishing functions, and seamlessly work with a variety of output devices and other software. Throw a network, voice mail, and image processing into the picture—all running under a GUI—and it's easy to

COMPANY INFORMATION

Advanced Logic Research, Inc.
9401 Jeronimo
Irvine, CA 92718
(714) 581-6770
Inquiry 872.

AST Research, Inc.
16215 Alton Pkwy.
Irvine, CA 92718
(714) 727-4141
Inquiry 873.

see a future where even the low person on the PC totem pole is going to require much more computing muscle than is common today.
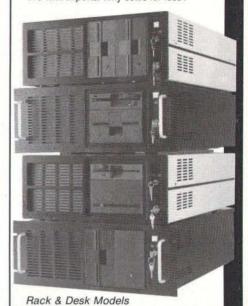
The handwriting is on the wall, but lesser systems still do the job and will continue to do so for the immediate future. In many cases, it's easier to justify purchasing a 386 system that can make use of a more powerful CPU than a machine that's stuck with one configuration. When the applications begin to burst out at the seams of the hardware, the upgrade will be relatively painless in terms of both cash spent and time lost due to setting up.

These systems are inexpensive only in the sense that you get a lot of processing power for the money. Neither the ALR PowerVEISA 486/33 nor the AST Premium 486/33 will be sales leaders, but they provide the power for those who need it now and, in lower-end configurations, a convenient way for others to upgrade as their circumstances and needs dictate. ∎

Michael Nadeau is BYTE's managing editor for reviews. You can reach him on BIX as "miken."

We've taken someth documents. The

Today, it takes more than powerful prose to make a document crisp, clear, and compelling.

It takes content. As well as style. In other words, charts, tables, figures, and formats previously beyond the reach of all but desktop publishing pros.

However, with new Microsoft® Word for Windows,™ you can turn out these documents with relative ease. Even if your strength is pronouns, not picas.

Word for Windows' intuitive graphical user interface and fully editable WYSIWYG make creating complex documents remarkably simple.

You can manipulate text, graphics and numbers more easily than ever before. And, you can see your changes as you make them, right on-screen.

Word for Windows even gives users a place to

Microsoft Word for Windows

ng out of complex
complex part.

start. A comprehensive, computer-based training program and context-sensitive, on-line help get you up and running quickly. And keep you there.

Of course we realize that you're not about to make the leap into graphical word processing if it means losing all of the work you have in your current program. Fair enough. Word for Windows allows you to directly read and write files from virtually all existing word processing programs.

Naturally, there's much more to Word for Windows than words can describe. So call (800) 541-1261, Dept. K47 and order our $9.95* Working Model. Once you've tried it, you'll discover that complex documents need only look that way.

Microsoft
Making it all make sense

Multiuser Databases:

SQL servers are not all created equal

Rich Finkelstein

According to conventional wisdom, all Structured Query Language (SQL) database servers are the same and have become "commodities," but nothing could be further from the truth. Each DBMS has distinct characteristics that make it either appropriate or inappropriate for a particular application. To further complicate matters, it seems that vendors place a high priority on adding features to their DBMS checklist—even though many of the features may offer little or no direct benefit. You should therefore look for the DBMS that best matches your requirements.

In spite of all the confusion, SQL servers have a lot to offer. Basically, a SQL server is a database engine that performs relational mathematics on sets of data. Unlike the case with a plain-vanilla relational DBMS, you don't have to tell a SQL server *how* to find the data you want; instead, you just ask for the data. The only catch is that you have to ask for the data using the server's SQL.

All the servers reviewed here support ANSI Level 1 SQL. Most support ANSI Level 2 and try to stay close to IBM's System Application Architecture (SAA) SQL. Some of the DBMS vendors have added useful extensions to their SQL, such as string, mathematical, and statistical functions, along with special data types, such as multigigabyte binary fields. Each SQL implementation differs considerably in this area, so you'll want to check for specific functions that are critical to your applications. Two SQL

servers in this review, Ingres and InterBase (see the text box "A New Kid in Town: InterBase" on page 142), let you define your own functions for added flexibility.

Each SQL server is equipped with a transaction manager that ensures that tables and indexes are synchronized—even after program or system failure. The problem is that when a program abruptly ends or the computer shuts down, a particular transaction might not have completed its updates, or data in the buffers might be lost. The transaction manager will detect this condition and automatically remove all partial updates. Therefore, tables and indexes reflect only committed (i.e., completed) transactions.

The SQL servers must also guard against data loss after a disk failure. They have backup/restore utilities that create and restore copies of a database. They also come equipped with forward recovery procedures that recover all updates that were completed between the last backup and the point of disk failure.

All SQL servers support *entity integrity* (i.e., each record in the table must be unique) via a combination of the SQL Unique Index facility and the Not Null column attribute. Another type of integrity, *referential integrity*, is best described by example. If you have two tables—say, a customer table and an invoice table—you must ensure that an invoice is not added to the database unless the customer exists in the customer file. Otherwise, the tables become out of sync, and thus the referential integrity will be lost. You also want similar integrity when you update and delete customer records.

Some SQL servers have no referential integrity. It is up to the user or programmer to handle the problem, with either program code or a series of interactive SQL commands. This will inevitably lead to unsynchronized tables. In some cases, depending on the application tool

used, it may not even be possible to enforce referential integrity from the application.
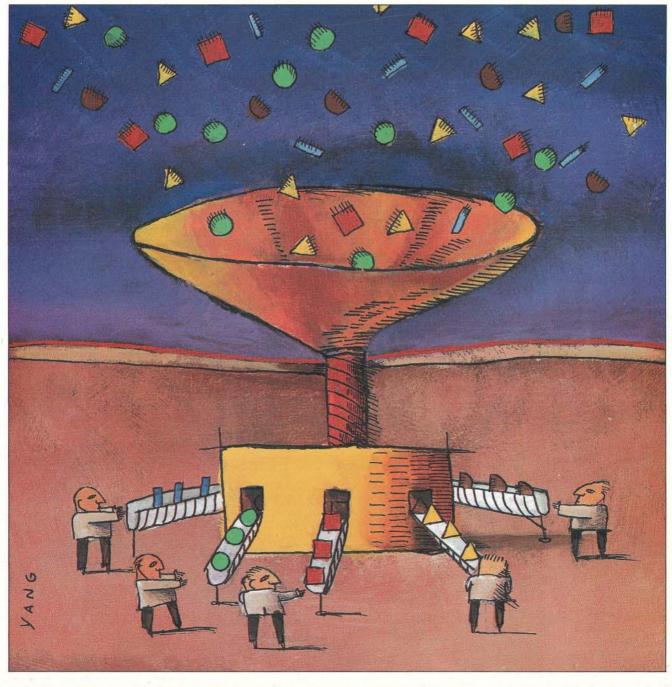
A much better method for handling referential integrity is by storing the referential rules in the DBMS "catalog," which is the body of commands and functions for that DBMS. Every time a table is updated, the referential integrity rules are automatically executed, or "triggered" (not to be confused with Oracle's SQL*Forms triggers, which are program-level code). This way, users and programmers no longer have to be concerned about referential integrity problems. Integrity rules are centralized so that they are consistent across all applications and are easier to maintain. Servers like SQL Server and Ingres support DBMS-level triggers that can be used to implement referential integrity as well as any other business rule.

The best referential integrity solution is one in which the DBMS vendor develops all the referential integrity logic and makes it available with the SQL CREATE TABLE or ALTER TABLE commands. Database administrators do not have to create any trigger logic, since the DBMS engine automatically maintains the referential integrity. As a result, there's much less of a chance of errors or omissions. IBM's OS/2 Extended Edition Database Manager (see the text box "A Preview of OS/2 Extended Edition" on page 146) implements this type of solution.

All the SQL servers except Progress have application programming interfaces (APIs) that support various languages such as C, COBOL, and Pascal. Progress implements its own 4GL language.

All the SQL servers reviewed here strive for high performance, especially if their makers sell their DBMSes in on-line transaction processing (OLTP) environments. Every server supports B-tree indexes for fast direct and sequential retrievals. They each also include query

The SQL

THE SQL FEATURES

The major features that set the SQL servers apart (● = yes; ○ = no).

SQL server	Price	Operating systems	4GL	Languages	Locking level	Query optimization	Index storage options	Data storage options	Referential integrity
XDB	$2590	DOS	Forms	C, COBOL, Pascal	Row	Cost-based	B-tree	None	Partial with DDL
SQLBase	$2495	DOS, OS/2	SQLWindows	C, COBOL	Page	Cost-based	B-tree	Hashed	None
SQL Server	$2495	OS/2, Unix, VAX/VMS	ABT Workbench	C	Page	Cost-based	B-tree, clustered	None	Triggers
Oracle	$2499	OS/2, Unix, VAX/VMS	SQL*Forms, SQL*Report	C, COBOL, FORTRAN	Row	Syntax-based	B-tree	Clustered tables	None
Ingres	$2490	Unix, VAX/VMS	ABF, QBF, RBF	C, COBOL, FORTRAN	Page	Cost-based	B-tree, ISAM	Hashed, sequential	Triggers
Informix	$1600–$10,000	Unix, VAX/VMS	Informix-4GL	C, COBOL	Row or page	Cost-based	B-tree, clustered	None	None
NetWare SQL	$795	NetWare VAP, NetWare 386, NLM	None	C, COBOL, Pascal	Row	Syntax-based	B-tree	None	None
Progress	$4500	DOS[2], Unix, VAX/VMS	Progress-4GL	None	Row	Cost-based	B-tree	None	Partial with file triggers

[1]No error checking or logic.   [2]Single-user version only.

optimizers that automatically select indexes when data is accessed or updated.

The servers all include automatic page or record locking that maximizes concurrency when many users are accessing the same date. Table-level locking would be the worst situation, since it would let only one user at a time access a table. Page-level locking locks several records at a time and promotes a higher level of concurrency. Record-level locking is the finest level of granularity for more concurrency.

All the servers also have automatic deadlock detection. Deadlocks occur when two or more transactions are waiting for records that are being held by other transactions. Since none of the transactions can continue processing, the SQL server must abort at least one transaction and send a message to the program to requeue the transaction.

These are the basic and necessary features of database servers. But each vendor has enhanced its server far beyond these basics.
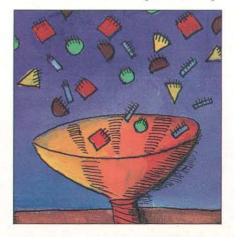
XDB

XDB is one of two SQL servers that can run under MS-DOS (an OS/2 version is also in the works). It will run on any NetBIOS LAN and requires only 640K bytes of memory. You can use a special 16-megabyte version to run in extended memory. Depending on your application, XDB should be able to handle up to 15 users with this additional memory.

XDB is highly compatible with IBM's mainframe DB2 DBMS. Most of the SQL DBMSes try to duplicate DB2's SQL syntax. XDB also duplicates DB2's messages and return/error codes. In addition, XDB displays data in the same way that DB2 would, and it maintains DB2's SQL restrictions.

XDB follows DB2's version 1.2 syntax closely. It includes the same data types, among them VARCHAR (variable characters) and LONG VARCHAR (long variable characters). XDB's LONG VARCHAR is currently limited to 1500 bytes, while DB2's can handle up to 3200 bytes. XDB also implements the CREATE SYNONYM statement, NOT NULL WITH DEFAULT, and CREATE VIEW WITH CHECK OPTION.

Release 2.3 of XDB copies DB2's op-

erators and functions. These include CURRENT TIMESTAMP, CURRENT TIME, CURRENT DATE, and CURRENT USER system variables. XDB uses DB2's || concatenation operator and supports several date and time formats, including those of the International Standards Organization, IBM USA, IBM European Standard, and Japanese Industrial Standard Christian.

These additions give XDB several levels of compatibility. XDB supports ANSI SQL Level 2, IBM's SAA, IBM's DB2 SQL, and its own XDB extensions. No other SQL DBMS can make this claim. XDB uses a compatibility flag to enforce a particular standard; if you want all users and programmers to use the DB2 standard and no other extensions or deviations, you can turn on the DB2 compatibility flag to permit only DB2 commands, functions, and data types.

These features make XDB the perfect platform for developing DB2 applications. Many large organizations use the Micro Focus or Realia COBOL compilers together with the Customer Information Control System option, which you can purchase in a package with XDB, to develop applications that will ultimately run under DB2.

You do not need NetWare to operate XDB on a LAN. Many people choose simpler network operating systems such as LANtastic, which supports printer and file sharing.

Stored procedures	BLOBs	User-defined functions	DB2 gateway	Optimized distributed query	Two-phase commit
○	○	○	○	○	○
Partial[1]	○	○	●	○	○
●	●	○	●	○	Partial
●	○	○	●	○	○
●	●	●	●	●	Partial
○	●	○	○	●	○
○	○	○	○	○	○
○	○	○	○	○	○

XDB on a server is exactly the same package as XDB on a single-user DOS PC. It is simple to install, and it has an interface that reflects its PC heritage. Many of the other servers have migrated from Unix or VAX systems and have user interfaces that are more appropriate for those environments than for PCs. They also carry excess baggage because they were developed for larger systems long before PCs were introduced.

In addition to its simplicity, XDB has many other attractive features. It has partial referential integrity in its SQL data definition language (DDL), and the company says that an upcoming version will extend its capabilities so that it will more closely emulate IBM's DB2. Currently, XDB supports PRIMARY KEY definition and INSERT referential integrity.

XDB implements row-level locking, which increases concurrency. XDB supports both of DB2's isolation levels: Cursor Stability and Repeatable Read. CS releases the row-level lock on the record after the program is finished processing it. RR maintains the locks until the transaction is completed, and it guarantees that the record will not be updated by other transactions until the transaction has ended.

XDB's weak spot is its application development tools. Its forms development system is showing its age and needs more functionality. You can use it to develop medium-size applications, but for more complex work, you will need to use the COBOL or C language interfaces. However, the company says that it is working on major enhancements to address many of these weaknesses.

SQLBase

The other DOS-based DBMS is SQLBase, one of three product lines from Gupta Technologies. The company also markets a windows application development tool called SQLWindows, and a mainframe cooperative processing link called SQLNetwork.

SQLBase was the first SQL database server for the PC. Initially it ran under DOS, but now it also runs under OS/2. (Gupta says that it will also have a Sun Unix version this year.) SQLBase will run on any NetBIOS network and also as a stand-alone DBMS under DOS.

SQLBase installation and administration procedures were designed with the PC user in mind. Its backup and recovery procedures exemplify this goal. You need only one command—BACKUP SNAPSHOT—to make an on-line backup of the database and its associated audit logs. You can use the RESTORE SNAPSHOT command to recover from disk failures. SQLBase automatically applies the correct backup and audit files without further user intervention.

Experienced administrators can use the BACKUP DATABASE/LOGS and RESTORE DATABASE/LOGS for finer control over backups and restores. After the databases and logs have been restored to the disk, the administrator has the option of recovering data to a specific point in time, to the point of disk failure, or to the time the on-line backup ended. The ROLLFORWARD TO END/TIME/BACKUP command supports these options.

SQLBase supports both B-tree indexes and hashed indexes (i.e., tables). B-tree indexes offer the best all-around performance, but if you need fast direct retrievals, then hashed tables can't be beat. Hashed tables let you access records directly without any indexes. A special randomizing algorithm inserts and retrieves records with only one I/O operation.

Since hashed table algorithms do not require any index retrievals, they eliminate at least one, and sometimes two or three, I/O operations for each record access. This can dramatically improve performance. You can also build B-tree indexes on any field in the record.

SQLBase's algorithm is especially good on tables with infrequent updates and serially numbered keys. You can define other types of tables to be hashed, but you must verify that the randomizer is doing a good job of randomizing the records. If the randomizing algorithm is not appropriate for the record's key, records will randomize to the same database page. Then when the page is filled, SQLBase chains that page to another page at the end of the database. This creates long synonym chains, and it's a poor use of space. Instead of saving I/O operations, the system ends up doing more.

SQLBase provides excellent support for mixing decision-support systems on the same server as OLTP applications. SQLBase's read consistency guarantees that a user who wants to only read a record will never be blocked by a transaction updating the same record. The update transactions will never interfere with queries and reports. Read consistency adds extra overhead, which hurts overall performance. The database administrator has the option of turning read consistency on or off.

SQLBase lets you use one command to store and execute a series of SQL commands. Gupta Technologies calls this facility "chained-SQL." Chained-SQL minimizes the interaction between the database server and the application program running on a PC workstation. Theoretically, chained-SQL can boost SQLBase's performance 40 percent or more; unfortunately, chained-SQL has limited practical usefulness.

continued

Chained-SQL does not let developers check for errors, nor does it allow any branch or control logic. Almost all applications require these capabilities, so it's unlikely that you will be able to take advantage of chained-SQL. Chained-SQL should be contrasted with SQL Server's and Ingres's stored procedures, which are fully functional and provide real performance advantages.

SQLBase supports COBOL and C language APIs. SQLBase's API permits backward fetches, which is important in PC applications.

SQLWindows is an excellent programmer's tool for developing sophisticated windows-based applications. SQLWindows has a long learning curve, but it's worth the effort if you need to build windows applications. With some refinement, SQLWindows can become the tool of choice for SQL database developers. Together with SQLNetwork, it can access DB2 databases. Gupta Technologies has announced that it is developing links to other DBMSes, including OS/2 Extended Edition Data Manager, Oracle, and SQL Server.

SQL Server
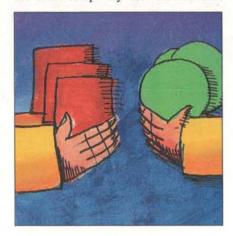
The Microsoft/Sybase SQL Server has long set the standard for client/server technology. SQL Server runs under various Unix systems, OS/2, and VAX/VMS. The OS/2 version of SQL Server requires named pipe network support, so you must install a version of LAN Manager (e.g., 3Com 3+Open, the Ungermann-Bass version, IBM LAN Server), or Novell's NetWare Requester for OS/2.

SQL Server has limited distributed update support (i.e., the ability to update several remote and local databases within a single transaction). It comes with function calls for coordinating updates across multiple databases, but it is the programmer's responsibility to issue the correct function calls. It is much better if the DBMS does all the work. SQL Server also supports remote procedure calls that allow transactions to execute procedures on other SQL Servers.

SQL Server supports referential integrity and other business rules with triggers. Triggers are small SQL programs, written in SQL Server's Transact-SQL language, that are stored in the DBMS catalog. Each trigger is associated with a particular table and a SQL update function (e.g., update, delete, and insert). They are automatically executed whenever a transaction updates the table. You can write triggers to enforce any database validation rule, including referential integrity.

Since they are stored in the catalog and automatically executed, triggers promote consistent integrity constraints across all transactions. The triggers are easy to maintain because they are stored in only one place—the DBMS catalog. Rules are enforced for any application that accesses the database, such as spreadsheet programs.

SQL Server also stores rules in its catalog. Rules apply to columns, and you use them to specify user-defined data types and range checks.

SQL Server implements a multithreaded, single-server architecture. This type of architecture is also used by SQLBase. Multithreaded servers perform most of their work and scheduling without interacting with the operating system. Instead of creating user processes, multithreaded servers create a thread for each new user. Threads are more efficient than processes, and they use less memory and CPU resources. In contrast, Oracle and XDB use a process/user architecture.

Its multithreaded architecture enables SQL Server to efficiently service a large number of users. It can service 40 users simultaneously on a 10-MB 33-MHz Compaq with only minor degradation in performance. However, SQL Server's single-server architecture does not allow it to take advantage of multiple processors. However, Sybase says it is working on a "virtual server" architecture that will create multiple servers on a single machine.

SQL Server was the first DBMS to implement stored procedures, which are similar to triggers. They are small Transact-SQL programs that are stored in the DBMS catalog. Any applications (e.g., databases and spreadsheets) can call a stored procedure. Instead of executing one SQL command at a time, stored procedures execute several commands simultaneously—without any further interaction with the application.

This saves a considerable amount of network overhead and can boost performance by 40 percent or more. Since Transact-SQL is a full language, developers can write complete procedures with branch-and-control logic, assignment operators, and error-checking capabilities. Oracle's OTEX and SQLBase's chained-SQL do not have these features.
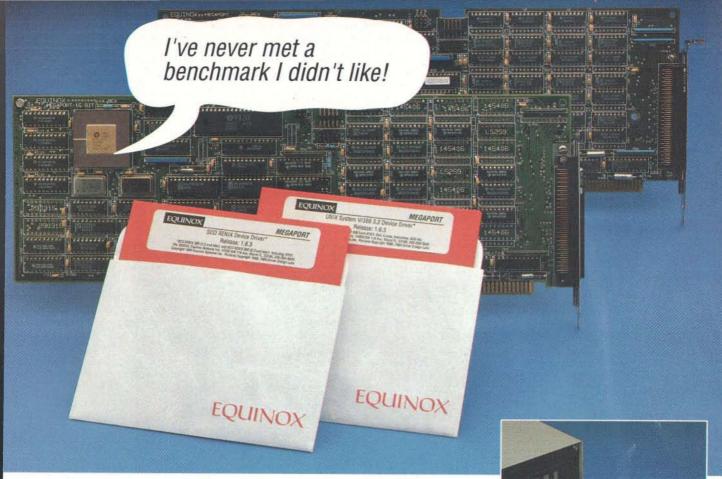
Clustered indexes is a little-known but highly valuable performance feature that greatly increases sequential processing performance. Clustered indexes maintain table data in a physical keyed sequence that enables sequential processing without the use of indexes.

Most other DBMSes must use indexes to sequentially retrieve a range of records or a whole table. This means that the transaction must perform at least one index I/O operation and one data I/O operation for each record. Often the DBMS must perform more. Suppose a transaction must retrieve 1000 records. It must then perform 2000 or more I/O operations.

Clustering reduces the number of I/O operations by eliminating the index I/O operations and clustering data in the same database page. If the database record size is 100 bytes and the database page size is 2000 bytes, then you can sequentially cluster 20 records to a page. This means that the transaction needs only 50 I/O operations (1000/20) to read all 1000 records. That's 40 times less I/O operations than required with simple B-trees. In my experience, this feature makes the critical difference in the success and failure of an OLTP application.

SQL Server has other useful features. Version 4.0, which is available on Unix platforms, supports 2-gigabyte binary fields that can store a lot of text or large images. It also has disk mirroring and fault-tolerance features. Sybase has also defined an Open Server specification that allows developers to build gateways between Sybase and other DBMSes such as DB2.

SQL Server has some weaknesses as well. It does not support the standard IBM SAA application "cursor" programming interface. A cursor stores the results of a SQL query and allows a program to move forward through the data one record at a time. Sometimes, as in the case of SQLBase, a programmer can move backward within a cursor. Without a cursor, it's harder to program transactions that must browse through data.

continued

A New Kid in Town: InterBase

Most PC users have never heard of Interbase Software, but that may change. Ashton-Tate recently bought a 51 percent equity interest in the company and plans to use Interbase Software technology in its next-generation SQL database server. Although for now the program InterBase does not run on a PC, there's little doubt that Ashton-Tate will exploit its advanced distributed database capabilities and its ability to handle what Interbase Software calls Online Complex Processing.

In the view of Interbase Software, OLCP environments are substantially different from on-line transaction processing. Typically, OLTP transactions tend to look alike and involve small chunks of high volume. OLCP environments are complex and unpredictable, involving database requests of arbitrary size. Examples of OLCP applications include manufacturing resource planning; real-time process control; medical, publishing, and financial trading; and CAD/CAM applications.

InterBase has many unique features that are oriented toward OLCP. Among those are a new type of trigger called an event trigger, which automatically executes a procedure whenever a particular event occurs. The event trigger is managed by the DBMS, not by the application. InterBase also supports data integrity triggers similar to those of Sybase and Ingres.

InterBase features BLOB (for binary large object) data types of unlimited size. Developers can write BLOB filters that standardize BLOB access and that hide BLOB characteristics from the program. Once a filter is written for a particular BLOB, you don't have to be concerned about how to handle it.

InterBase supports a peer-to-peer network architecture. In a peer-to-peer architecture, there is no central server, and InterBase allows any or all workstations to act as database servers. InterBase's distributed database capabilities include a two-phase commit that allows transactions to update multiple databases at multiple sites. Ingres and Sybase also support two-phase commits, but InterBase's is automatic and under DBMS control. Ingres and Sybase's implementations require programmer coordination.

OLCP transactions have long durations and lock large amounts of data. InterBase implements a special multiversioning locking technique that allows transactions to retrieve and hold records for long periods of time without blocking other transactions. This technique also works well for distributed OLCP transactions, but it's not clear how well it supports concurrent OLTP transactions. The same locking and storage mechanisms that make it appropriate for OLCP applications may make it inappropriate for OLTP.

Currently, InterBase runs on VAX/VMS, Apollo, and SunOS workstations. Ashton-Tate will very likely introduce a 386 version for OS/2 in the near future—possibly by the end of the year.

SQL Server does not have any application development toolkit under OS/2. Sybase says that it is porting its APT-Workbench toolkit to Presentation Manager, but that version will not be available until later this year. Until other products become available, developers must use SQL Server's C language API.

Oracle

Oracle Corp., the leading provider of Unix DBMSes, recently released its OS/2 database server. An Oracle database is portable to many different platforms as long as you maintain the same version of Oracle across all platforms.

Oracle claims that Oracle for OS/2 runs on a number of different networks, including Novell's IPX/SPX, NetBIOS, and TCP/IP. SQL*Connect allows users to connect to remote Oracle databases. Oracle's distributed database technology enables users to query local and remote Oracle databases within a single query. However, Oracle does not optimize distributed queries or support distributed updates. Oracle also has gateways to non-Oracle DBMSes such as DB2, but in practice these gateways have had stability and performance problems.

Oracle uses a process-per-user architecture. Each user connection demands its own server process. The up side of this architecture is that, within limits, it can use multiple processors. The downside is that it consumes memory resources and incurs extra CPU overhead. This can be a problem in single-processor 386 computers.

Oracle 6.0, the version that runs under OS/2, corrects many of the deficiencies of prior versions. Oracle 5.1, the most widely used version, locks full tables after an update. Oracle 6.0 now implements row-level locking using a time-stamping mechanism. Oracle supports a multiversioning mode that is similar to SQLBase's read consistency. This, together with its row-level locking, makes it especially suitable for mixed reporting and transaction environments.

However, Oracle still has a query and reporting weakness, and that is syntax-based optimization. This means that a SQL command behaves differently depending on the order of the table names in the SQL WHERE clause. This is a poor and antiquated solution. Instead of using a cost-based algorithm to determine the best way to join tables, Oracle puts the burden on users.

Skilled programmers may be able to find the right syntax, but if table sizes change, then you must modify the program. What's more, under certain circumstances, such as when you are using a third-party front-end product, the generated syntax may be completely out of programmer control.
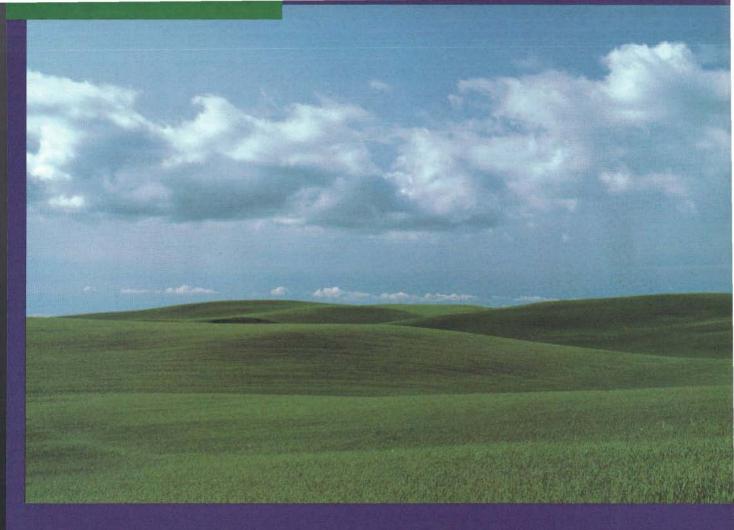
Although Oracle claims that its current version 6.0 performs much better than version 5.1, the claim is based on Oracle's OTEX feature. OTEX is similar to SQLBase's chained-SQL, and it can give Oracle a big performance boost. But OTEX has little practical value—it has no support for branch-and-control logic or any other programming logic. Like chained SQL, OTEX can only return the error code of the last SQL command executed.

SQL*Forms is Oracle's application development tool set. SQL*Forms has limitations that the company will address in version 3.0. This product is supposed to contain procedural capabilities in the form of Oracle's PL/SQL language, which the company is also working on.

Ingres

Ingres competes very closely with Oracle. It runs on a number of Unix and VAX

platforms. The manufacturer says that a version for OS/2 will be available by July. Other versions support the Sun386i and several networks, including Sun NFS, TCP/IP, and VAX DECnet.

Developers can access remote Ingres databases via Ingres/Net. You can use Ingres/Start for distributed updates and queries. Ingres optimizes distributed queries, but its distributed update capability, like SQL Server's, must be managed by the programmer. Ingres can access remote DBMSes, such as DB2, with its General Communications Architecture. In the past, Ingres's gateways have not been found to be reliable, but GCA appears to overcome these problems. GCA is the first gateway architecture that is based on the ANSI Remote Data Access standard.

Ingres is the only DBMS that supports a multithreaded, multiserver architecture. This overcomes the one-processor limitation of the multithreaded, single-server architectures. Administrators can start up multiple servers on a computer. These servers would be able to use multiple processors with no limitations. The advantage is that administrators would be able to direct transactions to specific servers, giving higher priority to critical transactions.
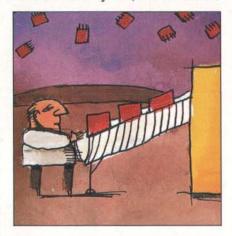
Ingres has many special features that make it technically superior to its competitors. For example, most other SQL DBMSes, with the exception of Oracle, use cost-based algorithms to optimize SQL queries. To determine the best index paths, they store row counts in their SQL catalog. Ingres stores not only counts but also histograms. Histograms more thoroughly describe data, since they capture its skewed nature.

Ingres also does a better job of interpreting SQL commands. There are many ways to express the same question within the SQL language. In most DBMSes, performance can be negatively affected by a nonoptimal formulation. For instance, some DBMSes perform better on joins, while others prefer subqueries. Ingres uses a "query-flattening" algorithm that ensures that all possible SQL constructs perform equally well.

Ingres supports hashed table structures as well as B-trees. As with SQL-Base, this gives the database administrator added physical storage flexibility. Ingres's Intelligent DBMS has recently been released on VAX and other high-end Unix platforms. The Knowledge Management Extensions (KME) includes stored procedures and triggers similar to those of SQL Server. Ingres's triggers and stored procedures are cre-

ated using Ingres's OSL language. Unlike SQL Server, Ingres allows multiple triggers per update operation per table. Procedures are therefore easier to write and maintain.

The KME contains a Resource Control System that establishes limits on resource consumption. With this facility, administrators can limit the number of rows returned by a query. This prevents runaway queries. Another facility, the Access Control System, allows database

administrators to control authority at the group level instead of the user level.

Skilled programmers can use the Object Management Extensions (OME) to define their own functions, operators, and data types. An example of a user-defined data type would be a geographical location. Instead of trying to describe geographic points and distances between points by using standard data types, skilled programmers can develop their own data type called POSITION. They can also create a user function called DISTANCE, which calculates the distance between POSITIONs. Users can then use these data types and functions directly in their SQL commands.

Ingres comes with a complete development tool set called Application-by-Forms. The company also offers an end-user query tool called Simplify.

The Santa Cruz Operation will package Ingres into its new 386-based Unix 5.2 operating system called Open Desktop. Under normal circumstances, Open Desktop with Ingres would be considered a good alternative to an OS/2 database server. Unfortunately, there seems to be very little awareness of the product. However, that could change.

Informix
Informix is very popular for small- and medium-size Unix-based applications. It is easier to administer than other Unix

DBMSes, it requires fewer hardware resources, and it offers excellent performance. Informix has great portability (supporting more than 250 platforms), as well as a good set of development tools that are oriented toward the professional programmer.

There are two versions of Informix. Informix-SE, Informix Software's original DBMS, runs on top of its C-ISAM file manager. This version runs under DOS, Unix, and DEC VMS. It does not have the performance characteristics of Informix-Online and is mainly used for applications with low transaction processing requirements. Access into Informix-SE is through a C-ISAM call interface or via SQL. Applications written with the SQL interface are compatible with the Informix-Online DBMS.

Informix-Online (formerly called Informix-Turbo) is Informix Software's entry into the OLTP marketplace. It does not run on as many platforms as Informix-SE, and it has not been ported to VAX/VMS or DOS, although an OS/2 port is in the works.

Informix-Star is Informix's distributed database system. Under Informix-Star, client workstations can access one or more remote Informix-Online databases. The distributed system uses Informix-Net to connect clients with the database servers.
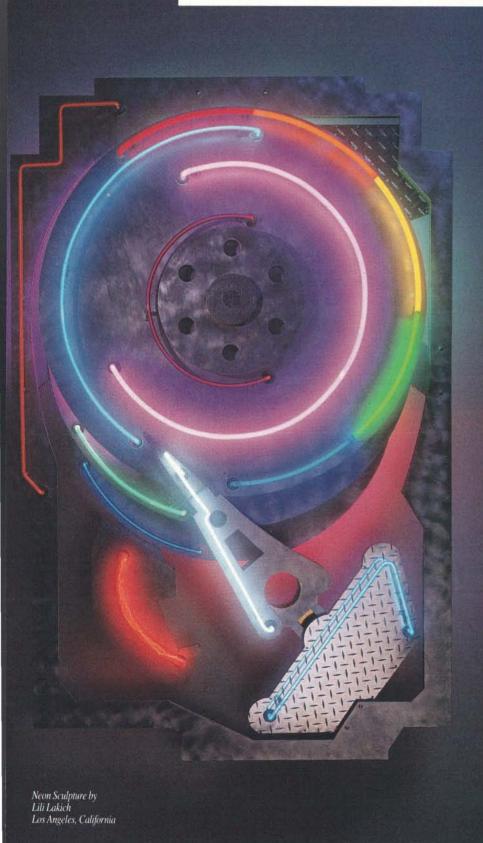
Informix-Net, which is the Informix network product, supports both TCP/IP and DECnet networks. Applications can read multiple databases but can update only one database at a time. A cost-based optimizer decides the best way to do a multitable join.

Informix uses a multiprocess shared-memory architecture. Each user requires his or her own process that can share memory with other processes. This is similar to Oracle's architecture. The advantage to this architecture is that it will make use of a multiprocessor computer. The disadvantage is that it requires more operating-system resources (i.e., memory and CPU) than the multithreaded architectures of Ingres and Sybase.

Informix has a flexible locking manager. Most DBMSes support only one type of locking—table, page, or record. Informix supports all three. Database administrators use the ALTER TABLE command to choose the locking level. Large tables can be assigned page-level locking, while smaller tables can have record-level locks. Locking levels can be changed dynamically from within a program.

Like SQL Server, Informix supports
continued

A Preview of OS/2 Extended Edition

IBM's OS/2 Extended Edition has two components: the Communications Manager (CM) and the Data Manager (DM). Version 1.1 of DM could be used only as a single-user SQL DBMS, since there was no way to access it on a network. Version 1.2 of OS/2 EE will come with a new facility called Remote Data Services (RDS), which will allow applications running on OS/2 or DOS workstations to access the DM database server.

RDS will route SQL calls to the database server using IBM's Advanced-Program-to-Program Communications (APPC) protocol. This protocol is supported by OS/2's CM. Applications using the DOS requester will access the server using NetBIOS.

Neither OS/2 nor DOS RDS requires a network operating system. But if you need remote printing or file sharing, you will need IBM's LAN Server NOS. LAN Server supports over 90 percent of LAN Manager's calls and is the only NOS that IBM certifies for OS/2 EE.

IBM also added referential integrity support to the SQL data definition language that closely follows DB2's imple-

mentation. The company has also completed its SQL implementation by adding the UNION operator.

OS/2 EE supports Cursor Stability as well as Repeatability Read isolation levels. It has a new isolation level called Uncommitted Read, which allows read-only transactions to bypass any update locks.

DM's Query Manager is vastly improved and makes good use of Presentation Manager's graphical interface. Developers and end users can create forms and reports using QM. You can develop more-sophisticated applications using the IBM language interfaces.

QM comes with COBOL/2, Pascal/2, and FORTRAN/2 language interfaces. It has an interface to a new procedures language that is based on IBM's mainframe REXX interpreter language. The procedures language can call and pass variables back and forth between QM's forms and reports. This is a fast way to build applications. A business graphics interface has also been added to QM that allows users to install a graphics package designed to work with the graphical interface.

DM includes an operational status tool that displays a snapshot of server activity, including database names, alias names, the time and date of the last backup, and how many applications are currently connected to a specific database.

DM's major weakness is that it has no utilities to recover from disk failures. This means that you have to re-enter all transactions that were entered between the time of the last backup and the disk failure. Only the SQL database server is missing this function. IBM says that it recognizes this as a problem and adds that it is working on a solution.

IBM also says that it is working on direct links between OS/2 EE and DB2 using CM's APPC protocol. Already, IBM has announced OS/2 Customer Information Control System CICS (CICS is IBM's mainframe teleprocessing monitor), which supports an APPC connection between a PC CICS program and a mainframe CICS program. IBM has announced plans for complete distributed database capabilities among all its database servers, but that's probably several years off.

clustered indexes, 2-gigabyte BLOB (for binary large objects) fields for long text and large images, and disk mirroring. It also has a database performance monitoring tool called DB-Monitor.

Simple applications can be developed using the PERFORM form-design system and the ACE report writer. More-complex applications will need Informix-4GL. Informix-4GL is halfway between a procedural and a nonprocedural language. Most systems can be developed entirely within Informix-4GL, but C and COBOL language interfaces are available. Informix-4GL/RDS is an interpretive version of Informix-4GL, which allows programmers to develop applications interactively and then execute them without compilation.

Informix Software has had some past failings in the area of user support. For example, users tend to complain about the company's slow response to fixing bugs and poor telephone support. Lately, however, the company has put extra effort into upgrading its support, and it appears that progress is being made in this area.

NetWare SQL
NetWare SQL has its origins in Softcraft's popular Btrieve file manager. Many corporations and value-added resellers (VARs) use Btrieve's indexed file manager together with procedural programming languages such as C, Pascal, and COBOL to build complex applications. A few years ago, Novell bought Softcraft and made it a division of the company. Btrieve has evolved into NetWare SQL and is now packaged with Novell's NetWare.

The latest version of NetWare SQL is a step up from Btrieve but is missing important implementation features, which causes it to lag behind its competitors. It still has several major shortcomings, such as the lack of a 4GL, stored procedures, hashed tables, clustered indexes, a good buffering algorithm, and good optimization.

NetWare SQL is the only SQL DBMS that runs as a Novell value-added process, which means that it does not require a separate server. Other SQL DBMSes need their own server, since they run under DOS or OS/2. Novell says

that NetWare SQL 386 will be written as a NetWare 386 Network Loadable Module (NLM). NetWare SQL has limited distributed capabilities. It can read records from a remote NetWare SQL server one record at a time using remote procedure calls.

The only major SQL feature that NetWare SQL is missing is support for null values. Forward recovery from disk failures is also missing, although this problem is supposed to be corrected in NetWare 386, which is due out sometime in the middle of the year.

NetWare SQL's transaction management system is not nearly as sophisticated as that of its competitors. It automatically locks tables unless programmers explicitly request record-level locks. Programmer-controlled locking adds to programming time and is susceptible to errors. In this day and age, that is simply unacceptable.

The soon-to-be-released NetWare SQL 386 will have some improved performance features—for example, a limited stored procedure capability and
continued

COMPANY INFORMATION

Gupta Technologies
(SQLBase)
1020 Marsh Rd., Suite 210
Menlo Park, CA 94025
(415) 321-9500
Inquiry 1071.

Informix Software, Inc.
(Informix)
4100 Bohannon Dr.
Menlo Park, CA 94025
(415) 322-4100
Inquiry 1072.

Ingres Corp.
(Ingres)
1080 Marina Village Pkwy.
Alameda, CA 94501
(415) 769-1400
Inquiry 1073.

Interbase Software Corp.
(InterBase)
209 Burlington Rd.
Bedford, MA 01730
(617) 271-0221
Inquiry 1074.

Microsoft Corp.
(SQL Server)
16011 Northeast 36th Way
Redmond, WA 98073
(206) 882-8080
Inquiry 1075.

Novell, Inc.
(NetWare SQL)
6034 West Courtyard Dr., Suite 220
Austin, TX 78730
(512) 346-8380
Inquiry 1076.

Oracle Corp.
(Oracle)
20 Davis Dr.
Belmont, CA 94002
(800) 672-2531
Inquiry 1077.

Progress Software Corp.
(Progress)
5 Oak Park
Bedford, MA 01730
(617) 275-4500
Inquiry 1078.

XDB Software, Inc.
(XDB)
7309 Baltimore Ave., Suite 219
College Park, MD 20740
(301) 779-5486
Inquiry 1079.

improved optimization. As an NLM, it should be able to take advantage of Net-Ware 386's considerably improved performance.

Initial reports on NetWare 386 are good, but it still has to be proven as a reliable platform for running mission-critical applications. There is some concern about NetWare 386's inability to protect programs from other NLMs. A runaway NLM can bring down a database server, which would be unacceptable for most application situations.

NetWare SQL 386 will also include an embedded SQL precompiler for C. This follows the ANSI and IBM proposed standards. The precompiler translates SQL commands into low-level call functions.

On one hand, I don't have anything against NetWare SQL, but on the other hand, it is difficult to find strong reasons to recommend it. The best reason to use NetWare SQL is its ability to run as an NLM. In the near future, NLM versions of XDB and SQLBase will be available, and that would eliminate NetWare SQL's one real advantage.

Progress
Progress is a favorite of VARs and small organizations. It has complete DBMS capabilities and an adequate 4GL. Progress's chief competitor is Informix, which runs on most of the same platforms as Progress. Like Informix, Progress is portable across dozens of Unix and VAX platforms. A single-user DOS version is also available. Progress operates over NetBIOS, TCP/IP, DECnet,

and Novell SPX/IPX networks.

Progress was one of the last true relational DBMSes to support SQL. Release 5.0 supports almost all of ANSI Level 2 SQL and most of IBM's SAA SQL. Progress does not have any procedural language interfaces, but it has embedded SQL support in its own 4GL. You can also use Progress interactively. Progress's embedded SQL uses a standard cursor mechanism that retrieves and manipulates data one record at a time. The technique is similar to the way other SQL servers embed SQL in C, COBOL, and Pascal.

Release 5 also implements a multithreaded, multiserver architecture that is similar to that of Ingres. It attempts to maximize the use of multiple processors while minimizing CPU and memory overhead. Progress has automatic record locking. Developers are given many methods of fine-tuning locking control

by defining the scope of a transaction. To do this, developers use the DO TRANS-ACTION and REPEAT TRANSAC-TION statements. Progress has a database monitor for managing multiuser performance, and a full set of data administration utilities.
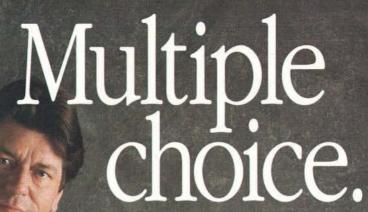
Progress has a data dictionary that stores field- and file-level validation rules. This is similar to SQL Server and Ingres's rules and triggers. In Progress's version, administrators use Progress's 4GL to define the rules. Unlike SQL Server and Ingres, which support file-level triggers for updates and deletes, Progress supports validations only on deletes—not on updates or inserts.

The heart of Progress is its 4GL. If you like the 4GL, you will like Progress. Since there is no other language interface, you must build all applications completely within the 4GL paradigm. I found it awkward and at times more difficult than procedural languages. Nonetheless, many developers swear by it, and if you can live within the constraints of a 4GL, you'll find you can be productive with Progress.

Perspectives on SQL
As you can see, each of the SQL servers offers unique characteristics that may make it ideal for a given application (see the features table on page 138). To put everything in a proper perspective, consider the following.

XDB is a good choice if you want a SQL server that is relatively inexpensive, runs on any NetBIOS network, and

continued

requires minimal data administration. It also makes sense if you're looking for a development platform for DB2 applications.

SQLBase does not have the portability of some of its competitors, such as Oracle, but it is easy to manage and has a good set of supporting tools. SQLBase will appeal to a wide range of developers and users who need a high-performance SQL database server with low resource requirements.

SQL Server is full of features, and it's an excellent choice for mission-critical OLTP applications. It is more costly in terms of resources and administration overhead than SQLBase and XDB, but for large networks and critical applications it can't be beat. Many corporations have even begun moving DB2 applications down to SQL Server.
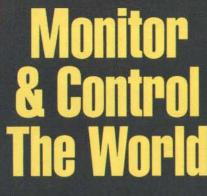
Oracle's strengths are its portability and large base of application support. Technologically, it lags behind its competitors, and it is a complex relational DBMS that is difficult to administer. It badly needs stored procedures and triggers. It should also implement either clustered indexes or hashed tables. The development tools also need upgrades.

Ingres's OME and KME extensions have established its technological leadership, which is a direction that other vendors should also consider.
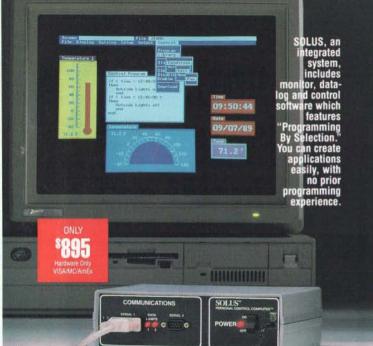
Informix is a perfect fit for smaller Unix systems. Informix's chief rival on Unix 386 platforms is the Ingres Open Desktop system. Ingres appears to have the edge in tools and administration facilities, while Informix has excellent VAR pricing and a large developer base. Both DBMSes are good picks, and customers will benefit from the competition.

Progress release 5 with SQL can now claim ANSI SQL compliance and has a rich 4GL to back it up. Unix developers enjoy Progress's aggressive pricing strategy that encourages VARs with low runtime costs. Prices are dependent on the type of platform and operating system but are very competitive with Informix's. This is especially true if you consider that the total Progress package is sold at one price, while Informix's Informix-4GL and Informix-RDS are priced separately. Progress provides a good alternative to Informix and Open Desktop—provided that you like its 4GL. ■

Rich Finkelstein, president of Performance Computing, Inc., in Chicago, is a leading consultant on relational and SQL-based database technology. You can reach him on BIX c/o "editors."

Circle 254 on Reader Service Card

25 Cartridges in One!™— 172 fonts and 20 symbol sets in one cartridge.

Not that long ago printing with more than one font seemed like it took 40 days and 40 nights. You had to wait to download soft fonts. Or you had to switch any number of cartridges in and out looking for the right font.

Then along came Pacific Data Products. Now you have fast and easy, yet economical access to more high-quality fonts and symbols than ever before. And there's also increased symbol set support including international symbols, currency and diacritical marks.

25 Cartridges in One! is ideal for almost any type of application—like word processing, presentation graphics, and spreadsheet publishing. With 19 printer drivers included, there's increased support for most of the popular software packages.

Now that it's loaded with 172 fonts and 20 symbol sets, go ahead—print up a storm.

25 Cartridges in One! is fully compatible with the HP LaserJet Series II, IID, and IIP. And, there's a new version available for the LaserJet III.

For the dealer that carries *25 Cartridges in One!* in your area, call Pacific Data Products at (619) 552-0880.

Pacific Data Products— Plug into Power

PACIFIC
DATA PRODUCTS

9125 Rehco Road, San Diego, California 92121
(619)552-0880, Fax (619)552-0889

FOR $1899, YOU CAN GET A LOT OF COMPUTERS.

World's Fastest Lunchbox

Dolch's P.A.C. 486-25
claims the portable
performance title, but
only for those with
very deep pockets

*Mark L. Van Name
and Bill Catchings*

I t was inevitable: 486 portables could not be far behind the first 486 desktop systems. The Dolch-P.A.C. (Portable Add-in Computer) 486-25 is the first—and, as we write this in late January, the only—486 portable.

The P.A.C. 486-25 is not a machine for the casual user. Its base system includes a 25-MHz i486 CPU, a socket for a Weitek WTL4167 floating-point coprocessor, 2 megabytes of memory, a 100-MB hard disk drive, a 5¼-inch 1.2-MB or 3½-inch 1.44-MB floppy disk drive, an electroluminescent CGA display, and three open AT-style expansion slots. Our evaluation unit came with the 3½-inch floppy disk drive, an additional 2 MB of memory ($780), a carrying case ($95), MS-DOS 3.3 ($150), and a VGA gas-plasma display instead of the standard EL screen ($595).

The P.A.C. 486-25 is also not for the poor user. The base system runs a hefty $12,995; our review unit would set you back $14,615. Dolch's 25-MHz P.A.C. 386-25 lunchbox portable in a comparable configuration costs $3100 less; you pay a huge premium for the i486.

Top Performance

Of course, you get something for all that extra money: performance. The P.A.C. 486-25 is by far the fastest portable computer that BYTE has ever tested. In the BYTE benchmarks, its overall application index of 26.3 is about 47 percent higher than that of the nearest competitor, the Toshiba T5200/100.

The i486 provides virtually all of that performance edge. The P.A.C. 486-25 beat the T5200/100 on BYTE's CPU index by about 57 percent. The i486's internal FPU increases the overall computing lead even farther; the system was more than three times faster on BYTE's FPU tests than the T5200/100, the Compaq Portable 386, or even Dolch's own P.A.C. 386-25.

The P.A.C. 486-25's video and disk subsystems, however, hold it back. The video system was slower than those of all three of the above-mentioned portables—a full 33 percent slower than that of the T5200/100. Several factors contributed

to this poor showing. First, the unit's Paradise VGA Plus adapter was an 8-bit card with only 256K bytes of memory. You can't replace the Paradise card with a faster one because Dolch modified the Paradise card to drive the gas-plasma display. The display actually connects to the card via the special display adapter on top of the card, not via the normal monitor connector. (There's also a special DC-to-DC converter that converts power from the unit's 200-watt supply into the power that the display requires.)

The disk subsystem's performance, while better than the display's, is also nothing to crow about. It turned in a BYTE disk I/O index that was 9 percent above the T5200/100's, but a full 16 percent below that of the P.A.C. 386-25. One reason for this mediocrity may be the overhead of translating the physical

continued

Dolch-P.A.C. 486-25

Company
Dolch Computer Systems
2029 O'Toole Ave.
San Jose, CA 95131
(800) 538-7506
(408) 435-1881

Components
Processor: 25-MHz Intel i486; socket for 25-MHz Weitek WTL4167 floating-point coprocessor
Memory: 2 MB of 80-ns DRAM in four 1-MB SIMMs on motherboard and 2 MB of RAM on 32-bit memory card (expandable to 16 MB); 128K bytes of BIOS ROM
Mass storage: 3½-inch 1.44-MB Toshiba floppy disk drive; 100-MB 23-ms IDE Conner Peripherals hard disk drive
Display: 11-inch gas-plasma VGA
Keyboard: 86-key AT layout with 12 function keys
I/O interfaces: One 9-pin and one 25-pin serial port; one 25-pin parallel port; five AT-style expansion slots (three open); one proprietary 32-bit memory card slot

Size
16 × 9½ × 7⅘ inches; 20 pounds

Price
Base system: $12,995
As reviewed: $14,615

Inquiry 861.

characteristics of the Conner Peripherals 100-MB IDE (integrated device electronics) hard disk drive into a configuration that the BIOS can handle. The overhead slows the drive, which offers a 23-millisecond average access time and a 1-to-1 interleave.

Despite this slow disk subsystem and display, the P.A.C. 486-25 is an amazing performer that made us want 486s on our desks.

Compatibility
The P.A.C. 486-25 fared wonderfully on our compatibility tests. Its standard AT-style expansion slots worked with all our test cards, including a Western Digital Ethernet controller. The system also had no trouble with a Microsoft Serial Mouse or the Xircom Pocket Ethernet Adapter.

On the software side, the unit flawlessly ran all our almost two dozen test programs, including several—Borland's Paradox 386 version 2.03, Foresight Resources' Drafix CAD Ultra 3.03C, Microsoft's Windows/386, Quarterdeck

Office Systems' QEMM-386 1.10, and Wolfram Research's Mathematica 1.2— that use 32-bit instructions and memory beyond 640K bytes.

Opening the Lunchbox
Despite its sterling PC compatibility and great performance, the P.A.C. 486-25 is stodgy-looking. Its almost 20-pound case is a dull gray rectangular box the color of an old Sears toolbox.

Inside, the P.A.C. 486-25 resembles a standard AT clone turned on its side. The vertically mounted Advanced Integration Research (AIR) motherboard runs along the front of the case, just behind the display. It measures 13 by 9 inches, the same size as the motherboard in many 386SX desktop systems, and sports over 75 chips—a lot, considering there's no memory or VGA support on it. There is only one application-specific IC, a Chips & Technologies integrated peripheral controller that manages the AT bus. The ROM BIOS is from Phoenix Technologies; the system by default copies both the ROM BIOS and the video BIOS to 32-bit memory for faster access.

The i486 in our evaluation unit was a B6 step chip, so it still had a few bugs in its FPU section. A Dolch spokesperson said that Intel plans to swap final i486s for B6 chips when the final chips are in full production; at that time, Dolch will in turn give early P.A.C. 486-25 buyers final i486s. The user can perform the swap (the CPU is socketed) or return the unit to Dolch at his or her expense.

The system's memory is on a separate AIR card in a proprietary 32-bit expansion slot. That card can hold up to 16 MB of memory in 16 1-MB single in-line memory modules. Our evaluation unit had 4 MB of 80-ns DRAM. The i486 contains a built-in 8K-byte cache and a cache controller that manages the system's memory and lets the CPU run without wait states about 85 percent of the time.
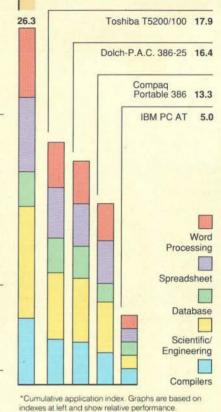
The P.A.C. 486-25 has five AT-style expansion slots: four 16-bit and one 8-bit. In our evaluation unit and in most P.A.C. 486-25 configurations, three of those slots are empty, providing a lot of expansion space for a portable. The other two slots hold the video and I/O cards. With a 5¼-inch floppy disk drive, one of the empty slots is limited to three-quarter length, but our unit's 3½-inch drive left all three slots clear to their full lengths. You reach the connectors on installed cards easily by removing a panel on the left side of the machine. To keep all the expansion cards and the i486 cool, *continued*

Dolch-P.A.C. 486-25

APPLICATION-LEVEL PERFORMANCE

Dolch-P.A.C 486-25 **26.3***

WORD PROCESSING

XyWrite III + 3.52 — Medium/Large

Load (large)	:11
Word count	:01/:06
Search/replace	:03/:10
End of document	:01/:07
Block move	:07/:07
Spelling check	:03/:18

Microsoft Word 4.0

Forward delete	:10

Aldus PageMaker 1.0a

Load document	:07
Change/bold	:12
Align right	:09
Cut 10 pages	:08
Place graphic	:02
Print to file	1:06

■ **Index:** 5.18

SPREADSHEET

Lotus 1-2-3 2.01

Block copy	:01
Recalc	:01
Load Monte Carlo	:05
Recalc Monte Carlo	:02
Load rlarge3	:01
Recalc rlarge3	:01
Recalc Goal-seek	:01

Microsoft Excel 2.0

Fill right	:03
Undo fill	:54
Recalc	:01
Load rlarge3	:10
Recalc rlarge3	:01

■ **Index:** 5.35

DATABASE

dBASE III + 1.1

Copy	:55
Index	:04
List	:52
Append	1:39
Delete	:02
Pack	1:17
Count	:03
Sort	:51

■ **Index:** 2.68

SCIENTIFIC/ENGINEERING

AutoCAD 2.52

Load SoftWest	1:12
Regen SoftWest	:16
Load StPauls	:05
Regen StPauls	:02
Hide/redraw	3:41

STATA 1.5

Graphics	:10
ANOVA	:06

MathCAD 2.0

IFS 800 pts.	:05
FFT/IFFT 1024 pts.	:05

■ **Index:** 8.15

COMPILERS

Microsoft C 5.0

XLisp compile	2:12

Turbo Pascal 4.0

Pascal S compile	:02

■ **Index:** 4.94

All times are in minutes:seconds. Indexes show relative performance; for all indexes, an 8-MHz IBM PC AT = 1.

26.3	Dolch-P.A.C 486-25
Toshiba T5200/100	**17.9**
Dolch-P.A.C. 386-25	**16.4**
Compaq Portable 386	**13.3**
IBM PC AT	**5.0**

Legend: Word Processing, Spreadsheet, Database, Scientific/Engineering, Compilers

*Cumulative application index. Graphs are based on indexes at left and show relative performance.

LOW-LEVEL PERFORMANCE[1]

Dolch-P.A.C. 486-25

CPU

Matrix	1.74
String Move	
Byte-wide	20.78
Word-wide:	
Odd-bnd.	17.04
Even-bnd.	12.31
Doubleword-wide:	
Odd-bnd.	13.55
Even-bnd.	8.07
Sieve	7.63
Sort	5.78

■ **Index:** 6.20

FLOATING POINT[2]

Math	1.64
Error	
Sine(x)	0.66
Error	
e^x	0.66
Error	

■ **Index:** 28.20

DISK I/O

Hard Seek[3]	
Outer track	3.26
Inner track	3.35
Half platter	6.70
Full platter	10.09
Average	5.85
DOS Seek	
1-sector	12.52
32-sector	22.06
File I/O[4]	
Seek	0.06
Read	0.48
Write	0.83
1-megabyte	
Write	3.39
Read	1.77

■ **Index:** 2.48

VIDEO

Text	
Mode 0	5.38
Mode 1	5.42
Mode 2	5.24
Mode 3	5.26
Mode 7	N/A
Graphics	
CGA:	
Mode 4	2.33
Mode 5	2.33
Mode 6	2.42
EGA:	
Mode 13	4.06
Mode 14	4.34
Mode 15	N/A
Mode 16	4.37
VGA:	
Mode 18	4.61
Mode 19	2.55
Hercules	N/A

■ **Index:** 1.81

N/A = Not applicable.

[1] All times are in seconds. Figures were generated using the 8088/8086 and 386 versions (1.1) of Small-C.
[2] The errors for Floating Point indicate the difference between expected and actual values, correct to 10 digits or rounded to 2 digits.
[3] Times reported by the Hard Seek and DOS Seek are for multiple seek operations (number of seeks performed currently set to 100).
[4] Read and write times for File I/O are in seconds per 64K bytes.
[5] For the Livermore Loops and Dhrystone tests only, higher numbers mean faster performance.

CONVENTIONAL BENCHMARKS

LINPACK	6420
Livermore Loops[5] (MFLOPS)	0.68
Dhrystone (MS C 5.0) (Dhry./sec.)	13192

Legend (graph): Toshiba T5200/100, Dolch-P.A.C. 386-25, Compaq Portable 386, IBM PC AT

Legend: CPU, FPU, Disk I/O, Video

For a full description of all the benchmarks, see "Introducing the New BYTE Benchmarks," June 1988 BYTE.

the P.A.C. 486-25 has two fans, one in the unit's upper right corner and one in the power supply.

Our only complaint about the P.A.C. 486-25's expansion abilities is that our evaluation unit's case was too snug. It bowed the cards slightly, so that getting them in and out was tricky.

Storage

Just below the upper fan is a drive assembly that contains the Conner hard disk

drive and a slim-line Toshiba floppy disk drive. A Chicony Electronics CH-101 I/O adapter controls the two drives. That card also provides two serial ports and one parallel port. If a 100-MB hard disk drive isn't big enough for you, Dolch also offers a 200-MB drive for $1995.

The BYTE Lab had some hard disk drive troubles: The first drive it received was bad, and the second produced intermittent read errors. We did not, however, experience those errors; the drive ran

without any problems throughout our tests. A Dolch spokesperson attributed the trouble with the first drive to damage during shipping.

The rest of the system is pretty much the same as other Dolch portables. The 86-key detachable keyboard clips onto the front of the unit. It follows the original AT keyboard layout, with two exceptions: There are 12 function keys, not 10, and those keys are in a row across the top of the keyboard rather than in a cluster on the left side. The keyboard had a mechanical click that was a little tinny for our tastes.

The only standard software was a set of two disks of Paradise VGA utilities and the built-in Phoenix setup program. Dolch sells MS-DOS, but you have to install it yourself.

The documentation is among the worst that we've seen. The manual was obviously assembled from different pages of other Dolch manuals—some pages were numbered sequentially, some by chapter—and it contained many small inaccuracies. A separate motherboard manual was somewhat helpful, but Dolch should put more time into its documentation. A Dolch spokesperson said that the company was in the process of preparing new manuals.
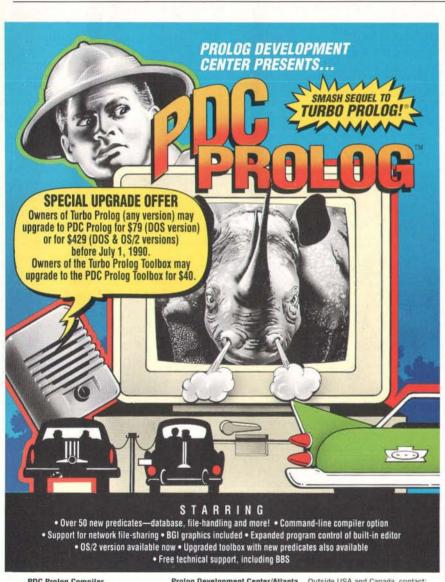
The Good and the Bad

On the support side, there's good news and bad news. The good news is a toll-free number for technical support and staff who know the system inside and out. The bad news is a standard one-year parts-and-labor warranty, under which you pay to ship the unit to Dolch for repair.

That warranty, like Dolch's software policies, was fine once, but not today. Vendors such as Dell, Austin, Gateway 2000, and PC Brand install DOS on the hard disk for you and provide a year of on-site service standard, even with their $2000 to $3000 clones. It's a shame that a company that's selling a system for $13,000 (or more) can't offer the same niceties.

The bottom line for the P.A.C. 486-25 is simple: It provides the best portable performance available, but at a very high price. If you need all that speed badly enough, and if you can afford the price, go for it. ∎

Mark L. Van Name and Bill Catchings are BYTE contributing editors. Both are also independent computer consultants and freelance writers based in Raleigh, North Carolina. You can reach them on BIX as "mvanname" and "wbc3," respectively.

Circle 201 on Reader Service Card (DEALERS: 202)

Four 386SXes to Go, Hold the AC

These 386SX laptops provide 386 power without the cord

David Claiborne

I t's no longer sufficient for a laptop to be small and lightweight and operate on batteries for several hours. As laptops have become more sophisticated, so have their users, who increasingly expect support for VGA displays and the processing power to accommodate power-hungry 386-based desktop applications. The problem is that, traditionally, these features have been mutually exclusive. To attain small size and weight and also achieve long battery life, processing speed and crisp displays have often been compromised.

The introduction of a new class of battery-powered 386SX laptop computers has changed all that. These computers take advantage of the 386SX CPU's lower power requirements to provide 386 processing power without being tethered to an AC power cord. The newest 386SX laptops all feature VGA graphics with a good aspect ratio, a minimum of 1 megabyte of memory, and a 20- or 40-MB hard disk drive in a 15-pound clamshell package that runs for 1½ to 3 hours on a single battery charge.

I recently tested four of these machines: the Dell 316LT, Tandon LT/386, Toshiba T3100SX, and Zenith SupersPort SX. The BYTE Lab benchmarks revealed some performance differences, but all these machines performed well, and all were significantly faster than comparable 286 laptops. They were also compatible with a variety of application programs, including Microsoft Word

Clockwise from top: The Tandon LT/386, Toshiba T3100SX, Zenith SupersPort SX, and Dell 316LT.

5.0, Lotus 1-2-3 release 3.0, and Borland's Paradox 386. There were, however, some definite performance differences and cost/performance trade-offs.

Dell 316LT
The 316LT is Dell Computer's first laptop, and it's impressive. The base system includes 1 MB of RAM, a 3½-inch 1.44-MB floppy disk drive, and a 20-MB 27-millisecond hard disk drive for $3499. The 40-MB 25-ms hard disk drive in my review machine boosted the price by

$300. Both hard disk drives are Conner Peripherals devices with an integrated device electronics (IDE) controller.

The 316LT sports attractive styling—something that's often absent in other machines. You'll find the on/off switch, reset button, and AC power jack on the left side of the unit, just under the keyboard. The battery pack slides in from the left side, just to the rear of the switches, and the floppy disk drive sits opposite the battery pack.

continued

	Dell 316LT	Tandon LT/386	Toshiba T3100SX
Company	Dell Computer Corp. 9505 Arboretum Blvd. Austin, TX 78759 (512) 338-4400	Tandon Computer Corp. 609 Science Dr. Moorpark, CA 93021 (805) 529-8227	Toshiba America Information Systems, Inc. Computer Systems Division 9740 Irvine Blvd. Irvine, CA 92718 (714) 583-3000
Components	**Processor:** 16-MHz 386SX **Memory:** 1 MB of RAM, expandable to 8 MB **Mass storage:** Internal 3½-inch 1.44-MB floppy disk drive; 40-MB 25-ms hard disk drive **Display:** 6- by 8-inch, 640- by 480-pixel VGA backlit LCD **Keyboard:** 83-key modified IBM Enhanced layout **I/O interfaces:** Serial, parallel, external analog VGA monitor, and keyboard ports; 8-bit half-card slot **Battery life (observed):** 1½ hours **Recharge time:** 2 hours	**Processor:** 16-MHz 386SX **Memory:** 1 MB of RAM, expandable to 5 MB **Mass storage:** Internal 3½-inch 1.44-MB floppy disk drive; 40-MB 25-ms hard disk drive **Display:** 6½- by 8½-inch, 640- by 480-pixel VGA backlit LCD **Keyboard:** 82-key modified IBM Enhanced layout **I/O interfaces:** Serial, parallel, and external analog VGA monitor, floppy disk drive, keyboard, and numeric keypad ports **Battery life (observed):** 3 hours **Recharge time:** 14 hours	**Processor:** 16-MHz 386SX **Memory:** 1 MB of RAM, expandable to 13 MB **Mass storage:** Internal 3½-inch 1.44-MB floppy disk drive; 40-MB 25-ms hard disk drive **Display:** 6- by 8-inch, 640- by 480-pixel VGA gas-plasma **Keyboard:** 86-key modified IBM Enhanced layout **I/O interfaces:** Serial (two), parallel, external analog VGA monitor, and numeric keypad ports; modem slot; proprietary expansion slot **Battery life (observed):** 3 hours (two batteries) **Recharge time (hours):** 3 hours (1½ hours per battery)
Size	3⅖ × 12⁷⁄₁₀ × 14⅗₁₀ inches; 15⅕ pounds	3⅖ × 12⅖ by 13½ inches; 14⅜ pounds	3¹⁄₁₀ × 12⅕ × 14⅕ inches; 14⁹⁄₁₀ pounds
Price	Base system: $3499 As reviewed: $3799	$4429	$5999
	Inquiry 867.	**Inquiry 868.**	**Inquiry 869.**

The external connections are at the rear, hidden behind a plastic panel that slides away to reveal connectors for a parallel port, a serial port, an analog VGA monitor, and an external keyboard. Also present are a half-length 8-bit expansion card slot and two external connectors for unannounced options.

The 83-key keyboard offers an arrangement that's similar to the IBM Enhanced layout, except that the Control and Alt keys only appear on the left side of the keyboard. Laptop users should experience no problems switching between the 316LT and their desktop computer.

The display is a backlit LCD that uses black characters on a white screen and measures 6 by 8 inches. A Chips & Technologies video BIOS provides full VGA capability.

Getting inside the 316LT is a simple matter of removing nine screws on the bottom of the case. The interior reveals three unique features. One is the half-length expansion card slot. Another is a set of eight single in-line memory module sockets that you can fill with 1-MB SIMMs for a total of 8 MB of memory. The third feature of note is the nickel-cadmium battery, which is the smallest

and lightest used by the machines reviewed: It weighs just 1⅔ pounds and measures 6 by 3⅕ by 1⅕ inches. (The Toshiba battery is actually lighter at 1½ pounds, but you need two of them.) It had a relatively short operating period of 90 minutes or less during testing. Most people will want a second battery. An internal backup battery powers the computer for up to 2 minutes, so you can swap in a freshly charged main battery without losing data.

You access Dell's ROM-based power management system by pressing the Ctrl/Alt/Insert keys. Options include shutting down the display and the hard disk drive during periods of extended inactivity, and disabling devices that you are not using, such as the serial and parallel ports, or a math coprocessor, if installed. Changes don't take effect until you reboot the computer.

The 316LT's bundled software consists of a tutorial and a few utilities—DOS isn't included. The utilities disk contains an extensive diagnostic program, an EMS 4.0–compatible expanded memory manager, and four utilities that adjust the display.

Since you can't take a Dell machine to

a dealer (the company sells exclusively through mail order), Dell takes extra pains to make sure its documentation will help its customers install upgrades and fix problems. The company provides complete and explicit instructions for taking the 316LT apart, adding memory SIMMs, adding an expansion board, and even swapping out the hard disk drive. But the manual is in a three-ring, slip-cased binder that's almost as large as the computer itself. Given that the 316LT is a portable computer, the manual should be portable, too.

Tandon LT/386

The Tandon LT/386 is the least attractive of the four units evaluated. The $4429 base machine includes 1 MB of RAM, a 3½-inch 1.44-MB floppy disk drive, and a 40-MB hard disk drive.

The unit's on/off switch is a standard rocker switch. Three external connectors (serial, parallel, and floppy disk drive port) are mounted on an exposed strip that's subject to damage during transit.

The 82-key keyboard layout is easy to use. It has Control and Alt keys on both sides of the keyboard and an Fn key at the lower left corner that invokes the Home,

PageDown, PageUp, and End functions. The 12 standard function keys are only half as high as the other keys, but I didn't find this to be a problem.

The Tandon display is the largest of the four, measuring 6½ by 8½ inches. The display, however, was one of the unit's weakest points. The hinge mechanism that lets you set the display to the optimum angle for the ambient light didn't work. Tandon says that this is not a problem with currently shipping machines. The screen is noticeably slower than the other LCDs; when I rapidly moved a mouse in Microsoft Word 5.0, the arrow indicator disappeared from the screen. But one nice feature of the display is the ability to toggle between black on white, white on black, and an external monitor.

Although Tandon, Dell, and Zenith all use the Conner Peripherals 40-MB hard disk drive, only Tandon optimized its BIOS to work with it. The resulting improvement in the drive's transfer rate didn't show up on BYTE's low-level disk benchmarks. But the LT/386 achieved impressive scores in disk-intensive application tests such as the dBASE III tests. In fact, the enhanced performance of the hard disk drive gave the Tandon the high-

est application index of the group.

Unlike Dell, Tandon doesn't want the user to open its computer. A proprietary slot is accessible by simply removing an external panel and sliding in a card. Tandon is developing a modem card and a card to link an expansion chassis to the LT/386. Neither was available at press time.

The Tandon system board accepts an 80387SX math coprocessor and an additional 4 MB of memory on two special memory cards. The instructions for adding a math coprocessor or extra memory are simple: "Contact your Tandon dealer for assistance." Despite this lack of direction, it is easy to open the unit and locate the empty socket for the coprocessor. But there's no mention of how to install the memory cards. Tandon says that it will include instructions in future revisions of the user's manual.

The Tandon battery is another weak link and demonstrates less engineering development than other laptop batteries. The nickel-cadmium battery floats loosely inside its storage cavity and rattles around when you move the computer. Tandon rates battery life at 3 hours, assuming no accessories and a 10 percent disk duty cycle. To conserve power, you can set the display and disk drive to turn off after a user-selectable time-out period ranging from a few seconds to 15 minutes. As with the Dell 316LT, the power management functions are in ROM. You access them by pressing Ctrl/Alt/Esc. The machine reboots itself upon exiting the setup program.

Like the other units, the LT/386 uses an external AC transformer, about the size of a small brick, to provide AC power and to recharge the battery. But where the other units recharge the battery pack in 2 to 3 hours, the LT/386 takes up to 14 hours. The time is even longer if you're using the computer during the recharge period.

The LT/386 comes with Tandon's version of MS-DOS 3.3. The standard software includes the Microsoft expanded-memory device driver for 386-equipped computers. There are no diagnostics or tutorial programs.

Documentation consists of a thin, wire-bound manual that covers both the LT/386 and the LT/286, and a DOS manual. The manual provides information on setting up the computer, a short overview of the DOS command structure, and a cursory guide to installing options.

Toshiba T3100SX

Toshiba's 386SX laptop breaks ground in several areas, maintaining Toshiba's

position as a major innovator in the laptop field. The machine's biggest achievement is the low-power gas-plasma display that makes the technology feasible in a battery-operated laptop. Even with the low-power design, Toshiba had to make concessions elsewhere in the machine to obtain sufficient power for the display, but the screen is very sharp.

The basic configuration includes a 3½-inch 1.44-MB floppy disk drive, a 40-MB hard disk drive, and 1 MB of RAM. But the T3100SX can accept up to 13 MB of internal memory. The unit also includes dual nickel-cadmium batteries that power the machine for up to 3 hours and let you swap out a depleted battery with the power on.

The on/off switch is located on the right side, just behind the floppy disk drive. Located on the left side are openings for an optional internal modem, an optional numeric keypad, and a reset button. In the rear of the machine, the remaining external ports are tucked behind the carrying handle. These include two serial ports, a parallel port, and an RGB port for an external VGA monitor. The parallel port accommodates a printer or an external floppy disk drive. Toshiba also includes a proprietary 16-bit expansion slot.

The 86-key IBM Enhanced–style keyboard is well laid out. The Fn key toggles the computer speed, displays the battery monitor window, and invokes the numeric keypad overlay.

The gas-plasma screen is impressive. It measures 6 by 8 inches. This provides an ideal 4-to-3 aspect ratio for VGA. It's sharp enough and fast enough to use as your desktop display. The internal circuitry maps the VGA color palettes into 16 intensity levels. Using VHCAD, a Toshiba utility program, you can assign VGA colors to specific gray levels. The VHCAD program includes 10 preset combinations that have the gray scales already optimized for most applications. The screen is the slowest of the displays tested, but what it lacks in speed, it more than makes up for in clarity.

Toshiba compromised hard disk drive performance to get the power for the plasma display. The T3100SX uses a specially built hard disk drive with a slower than normal rotation speed to save power. The slower rotation speed doesn't slow the average access speed—the drive still performs at around 25 ms. But it does slow the transfer rate in BYTE's benchmark tests. In the 1-MB write/read tests, the hard disk drive is substantially slower than those of the other units.

continued

Tandon LT/386, Zenith SupersPort SX, Dell 316LT, Toshiba T3100SX

APPLICATION-LEVEL PERFORMANCE

WORD PROCESSING XyWrite III+ 3.52	Tandon Med./Large	Zenith Med./Large	Dell Med./Large	Toshiba Med/Large
Load (large)	:10	:13	:14	:19
Word count	:04/:28	:04/:27	:04/:27	:04/:28
Search/replace	:07/:26	:07/:26	:07/:27	:07/:28
End of document	:02/:15	:02/:15	:02/:15	:02/:18
Block move	:16/:16	:14/:13	:14/:14	:11/:10
Spelling check	:11/1:24	:11/1:23	:11/1:24	:12/1:28
Microsoft Word 4.0				
Forward delete	:20	:19	:19	:19
Aldus PageMaker 1.0a				
Load document	:09	:09	:09	:13
Change/bold	:32	:33	:27	:37
Align right	:24	:24	:26	:28
Cut 10 pages	:23	:20	:30	:31
Place graphic	:06	:04	:05	:06
Print to file	2:02	2:14	2:16	2:43
Index:	**2.10**	**2.17**	**2.09**	**1.94**

SPREADSHEET Lotus 1-2-3 2.01	Tandon	Zenith	Dell	Toshiba
Block copy	:05	:05	:04	:05
Recalc	:02	:02	:02	:02
Load Monte Carlo	:18	:18	:18	:27
Recalc Monte Carlo	:09	:09	:09	:09
Load rlarge3	:05	:04	:05	:06
Recalc rlarge3	:02	:01	:02	:02
Recalc Goal-seek	:05	:05	:05	:05
Microsoft Excel 2.0				
Fill right	:07	:06	:06	:07
Undo fill	2:51	2:08	2:06	2:41
Recalc	:01	:02	:02	:01
Load rlarge3	:31	:25	:29	:35
Recalc rlarge3	:02	:02	:02	:02
Index:	**1.96**	**2.14**	**1.97**	**1.91**

DATABASE dBASE III+ 1.1	Tandon	Zenith	Dell	Toshiba
Copy	:36	:55	1:03	1:17
Index	:06	:06	:06	:29
List	1:33	1:11	1:09	1:34
Append	1:23	2:03	2:05	2:38
Delete	:02	:03	:04	:05
Pack	:55	1:35	1:36	2:09
Count	:04	:05	:05	:26
Sort	:33	:57	:58	1:52
Index:	**2.72**	**2.05**	**1.94**	**1.02**

SCIENTIFIC/ENGINEERING AutoCAD 2.52	Tandon	Zenith	Dell	Toshiba
Load SoftWest	2:36	2:49	2:42	2:58
Regen SoftWest	2:24	2:39	2:31	2:53
Load StPauls	:46	:46	:46	:51
Regen StPauls	:41	:41	:42	:44
Hide/redraw	35:21	36:56	36:31	37:15
STATA 1.5				
Graphics	1:16	1:11	1:11	1:11
ANOVA	:49	:45	:46	:47
MathCAD 2.0				
IFS 800 pts.	1:27	1:27	1:30	1:35
FFT/IFFT 1024 pts.	1:45	1:44	1:48	1:53
Index:	**0.84**	**0.83**	**0.83**	**0.78**

COMPILERS Microsoft C 5.0	Tandon	Zenith	Dell	Toshiba
XLisp compile	4:38	4:43	5:03	5:41
Turbo Pascal 4.0				
Pascal S compile	:05	:05	:05	:07
Index:	**2.17**	**2.13**	**2.13**	**1.62**

All times are in minutes:seconds. Indexes show relative performance; for all indexes, an 8-MHz IBM PC AT=1.

LOW-LEVEL PERFORMANCE[1]

CPU	Tandon	Zenith	Dell	Toshiba
Matrix	6.66	6.61	7.01	7.58
String Move				
Byte-wide	51.46	50.79	52.49	39.82
Word-wide:				
Odd-bnd.	42.73	42.13	46.23	44.71
Even-bnd.	25.76	25.40	26.27	19.94
Doubleword-wide:				
Odd-bnd.	27.79	27.44	29.77	29.93
Even-bnd.	19.44	17.24	21.37	19.88
Sieve	33.32	32.79	37.13	44.61
Sort	31.86	31.43	31.58	35.85
Index:	**1.90**	**1.93**	**1.82**	**1.88**

FLOATING POINT[2]	Tandon	Zenith	Dell	Toshiba
Math	N/A	N/A	N/A	N/A
Error				
Sine(x)	N/A	N/A	N/A	N/A
Error				
e^x	N/A	N/A	N/A	N/A
Error				
Index:	**N/A**	**N/A**	**N/A**	**N/A**

DISK I/O	Tandon	Zenith	Dell	Toshiba
Hard Seek[3]				
Outer track	3.34	3.32	3.37	5.08
Inner track	8.36	3.37	3.39	5.11
Half platter	8.42	8.39	8.02	10.29
Full platter	11.79	10.02	10.98	10.29
Average	7.98	6.28	6.44	7.69
DOS Seek				
1-sector	14.97	12.71	12.81	12.26
32-sector	37.66	23.70	24.76	36.96
File I/O[4]				
Seek	0.18	0.07	0.10	0.05
Read	0.55	0.56	0.56	1.44
Write	0.84	1.00	0.92	1.36
1-megabyte				
Write	3.19	3.59	3.29	5.20
Read	3.06	3.06	3.08	7.06
Index:	**1.76**	**2.12**	**2.05**	**1.50**

VIDEO Text	Tandon	Zenith	Dell	Toshiba
Mode 0	5.95	5.11	5.09	5.64
Mode 1	5.93	5.11	5.06	5.64
Mode 2	6.59	5.24	5.27	5.59
Mode 3	6.57	5.25	5.27	5.58
Mode 7	N/A	N/A	N/A	N/A
Graphics				
CGA:				
Mode 4	2.49	2.44	2.60	2.67
Mode 5	2.53	2.41	2.60	2.68
Mode 6	2.80	2.71	2.86	2.86
EGA:				
Mode 13	4.56	3.86	4.06	4.74
Mode 14	5.22	4.50	4.67	5.17
Mode 15	N/A	N/A	N/A	N/A
Mode 16	5.22	4.46	4.67	5.17
VGA:				
Mode 18	5.44	4.69	4.89	5.40
Mode 19	2.76	2.91	2.86	2.91
Hercules	N/A	N/A	N/A	N/A
Index:	**1.56**	**1.81**	**1.77**	**1.63**

N/A = Not applicable.

[1] All times are in seconds. Figures were generated using the 8088/8086 and 386 versions (1.1) of Small-C.

[2] The errors for Floating Point indicate the difference between expected and actual values, correct to 10 digits or rounded to 2 digits.

[3] Times reported by the Hard Seek and DOS Seek are for multiple seek operations (number of seeks performed currently set to 100).

[4] Read and write times for File I/O are in seconds per 64K bytes.

[5] For the Livermore Loops and Dhrystone tests only, higher numbers mean faster performance.

CONVENTIONAL BENCHMARKS	Tandon	Zenith	Dell	Toshiba
LINPACK	2942.79	2901.44	2991.52	3123.28
Livermore Loops[5] (MFLOPS)	0.0069	0.0069	0.0072	0.0068
Dhrystone (MS C 5.0) (Dhry./sec.)	3475	3840	3501	3410

For a full description of all the benchmarks, see "Introducing the New BYTE Benchmarks," June 1988 BYTE.

Tandon LT/386 **9.8***

Zenith SupersPort SX **9.3**

Dell 316LT **9.0**

Toshiba T3100 SX **7.3**

IBM PC AT **5.0**

9.8 9.3 9.0

7.3

- Word Processing
- Spreadsheet
- Database
- Scientific/ Engineering
- Compilers

*Cumulative application index. Graphs are based on indexes at left and show relative performance.

Tandon LT/386

Zenith SupersPort SX

Dell 316LT

Toshiba T3100 SX

IBM PC AT

- CPU
- FPU
- Disk I/O
- Video
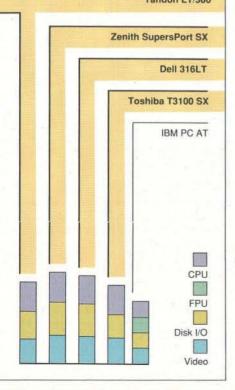
Toshiba's laptop experience is most evident in its power management system. First there are the dual battery packs. When one battery is depleted, the unit automatically switches to the second battery. Built-in charging circuitry indicates when a battery is being charged or is fully charged. The usual device power time-out and deactivation controls are activated from a ROM-based setup program that implements changes without requiring a system reboot. The auto-resume feature lets you turn off the computer while running an application and then resume at the same point when you turn the machine on again. Each battery is good for about 90 minutes, yielding a total of 3 hours of running time. Recharge time is about 3 hours for both batteries, assuming that the machine is turned off.

Toshiba includes an extensive amount of software with the T3100SX. Microsoft DOS 4.01 is standard, as are Quarterdeck's QEMM, an expanded-memory manager for the 386SX, and PC-Kwik Power Pak, which includes a disk cache, a print spooler, a RAM disk, and screen and keyboard accelerators. Also included is an on-line hypertext manual that covers both MS-DOS 4.01 and the T3100SX itself.

Written documentation includes a detailed manual that makes use of screen images, color photographs, and line drawings; two DOS manuals; and two small operation guides.

While Dell and Tandon were still planning accessories at press time, Toshiba supplies a wide range of options for the T3100SX, including an external battery charger ($399) that holds up to three battery packs, an external keypad ($99), an expansion chassis with five IBM PC–compatible slots ($999), and an Ethernet card ($699).

Zenith SupersPort SX

Zenith, a longtime player in the laptop field, has also developed a steady stream of innovative laptops. The SupersPort SX carries on the tradition. The unit borrows a great deal from its predecessors, but it improves performance substantially with the 386SX processor. The base configuration and price are the same as the Toshiba's.

The SupersPort SX uses the standard Zenith battery pack that attaches to the rear of the unit. This gives the SupersPort the ability to operate without its battery. If you're going to be near a power source, you can leave the battery pack at home, shedding 5 pounds and 3

continued

inches of depth from the machine's dimensions. When connected to the AC adapter/charger, the battery recharges in 3 hours. Extra batteries are $289 each.

The on/off switch, external keyboard socket, and floppy disk drive are located on the right side of the unit. The left side has an opening for a 1200-bps ($299) or 2400-bps ($499) internal modem. A sliding plastic panel at the rear of the unit protects a serial port, a parallel port, an analog VGA monitor port, an expansion chassis connector (a three-slot expansion box is $499), and an external floppy disk drive interface. Only the serial and parallel ports are accessible when the battery pack is attached.

The 79-key keyboard is the weakest part of the SuperSport SX. The key arrangement makes the transition from laptop to desktop difficult. The most serious offenders are the Escape key, located below the F1 key rather than in the top left corner, the backslash key, located to the right of the Enter key, and the Insert and Delete keys, located in the top right corner. To use F11 and F12, you must press Fn/F1 and Fn/F2, respectively. The SuperSport's display is a fluorescent backlit, VGA-compatible LCD with 640-by-480-pixel resolution and a 6- by 8-inch screen.

The SuperSport's display is noticeably faster than the Tandon's, but it isn't as clear as the Dell's. Like Tandon, Zenith doesn't want users opening up its machine. The company provides access to the math coprocessor, ROM BIOS, backup battery, and modem via a series of hatches and access panels on the bottom of the computer. Installing additional memory, however, requires removing the cover. Here Zenith provides no instructions, except to refer you to an "approved service center."

The unit's 5-pound battery provides up to 3 hours of operation, provided you use power judiciously. Power management on the SuperSport lets you shut down various system components when they aren't needed. The setup program is part of a ROM-resident monitor program, invoked by pressing Ctrl/Alt/Insert. The setup program features three screens, each containing a different configuration for battery operation, AC operation, and AC operation with an expansion chassis. The setup program can disable the modem, change the computer operating speed, and set time-out periods for the display and hard disk drive. The system reboots whenever changes are made.

Zenith includes MS-DOS 3.3. The ROM-resident monitor program contains both setup and diagnostics functions. The documentation is sparse. Two DOS manuals are complete, but the documentation supplied on the SuperSport itself provides details on only the externals of the computer.

The Right Choice

All four machines present price/performance trade-offs. Choosing the right machine depends on your needs. In terms of cost, the computers are grouped much more tightly than the list prices indicate. Dealer discounts usually shave 20 percent to 30 percent or more off the Zenith, Toshiba, and Tandon units.
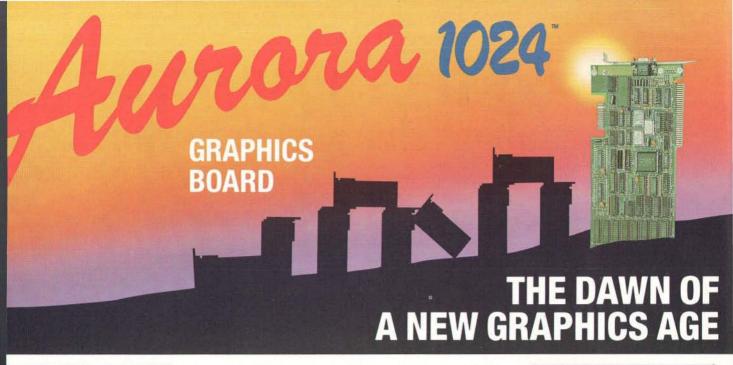
Performance-wise, no machine has the best of everything. The Dell 316LT offers the easy memory expansion and a crisp paper-white LCD, but its performance was middle-of-the-road, and its battery life should be longer. The Tandon LT/386 took top honors on the application index, due in large part to the close integration of its hard disk drive, but its overall construction leaves much to be desired. Also, both Dell and Tandon are planning, but didn't yet have, internal modems available. Toshiba and Zenith have modems available for their laptops.
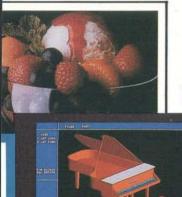
If detailed graphics are a requirement, the Toshiba T3100SX eliminates the competition with its gas-plasma display. The graphics come at the cost of a slower disk drive, as well as the slower video performance that gave the system the lowest application index of the four machines. But the T3100SX has many innovations, particularly in the area of power management, to recommend it as the top choice.

The Zenith SuperSport SX achieved the highest marks in terms of raw CPU, disk I/O, and video performance. The SuperSport suffers, however, from a nontraditional keyboard arrangement, and, at 16⁷⁄₁₀ pounds with the battery attached, it's almost 2 pounds heavier than the Toshiba.

The Toshiba T3100SX is my choice. The display is delightful to work with. If you're going to use CAD applications, desktop publishing, or anything else that requires a crisp display, the Toshiba is the top pick. If top screen quality is not a hard requirement, I'd choose the Dell 316LT. It has several unique design features, and it provides above-average performance at below-average price. ∎

David Claiborne is a computer consultant and freelance writer based in Highland, Maryland. He can be reached on BIX c/o "editors."

AGI Computers....
Unbeatable Value and Performance

386-25MHz Desktop

5.9* $4500 5.2* $5695

AGI	AST
3000L	Premium
4MB	2MB

■ MIPS
■ List Price

386-33MHz Desktop

7.9* $5100 6.7* $6595

AGI	AST
3000K	Premium
4MB	2MB

■ MIPS
■ List Price

486-25MHz Upright

11.2* $7995 10.95* $8995

AGI	AST
3900D	486/25t
4MB	2MB

■ MIPS
■ List Price

Sources: * Power Meter MIPs Version 1.5, The Database Group, Inc.

AGI -- Advanced Group Innovations.

The name that is synonymous with **reliability**. We have shipped over one half of a million computers in a short period of time. Now, we are excited to offer our new line of products for the '90s.

Try our 80486 based newcomer, or our overwhelmingly popular 80386 systems. The price to performance ratio of AGI machines leaves others way behind !

For a single or multi-user system, you have a choice of a laptop, a desktop, or even an upright.

Whichever you choose, AGI computer is the name you can unequivocally count on. Call us today.

For Information, Call
(415) 683-2800
Fax: (415) 683-4735
48460 Kato Road, Fremont, CA 94538

AGI 286/12
Laptop

AGI Computer, Inc.

Advanced Group Innovations

AGI is a trademark of AGI Computer Inc. AST is a trademark of AST Research Inc. AGIad2v.1 3/90

Circle 11 on Reader Service Card

Power and the Single User

Opus mingles the worlds of RISC and PCs with delightful results

Tom Yager

Power. The lust for it is pervasive not only on Capitol Hill, but on the desks of users and developers as well. Ever since the introduction of the first PC, the word heard most often from computer users has been *more*.

The truly power-hungry have a friend in Opus Systems, which has introduced the Personal Mainframes. These machines consist of a Motorola 88000 RISC coprocessor card sitting in the expansion slot of a 16-bit IBM AT clone. PMs (eight models were available at the time of this writing) integrate easily into the PC world, with the most unusual construction I've ever seen in Unix machines.

What really surprised me is that these are relatively inexpensive machines—the portable version is base priced at $13,995. [Editor's note: *This low-cost portable system, the 8110-PM, was evaluated in a January Short Take and was later obtained for benchmarking. The benchmark table on page 171 lists the test results for a 20-MHz system with 8 megabytes of memory.*] They turned in BYTE benchmark figures significantly higher than those of systems costing much more, delivering what appears to be the best performance for the money.

Unix Veteran

Opus has been in the Unix business for quite a while and has always emphasized both performance and ease of integration. As for the latter, Opus may have the market cornered. Everything that it sells

At $18,995, the Opus 8140-PM tower gives software developers and workstation users a bargain in high-end performance.

starts with an Industry Standard Architecture–bus coprocessor card. These cards offer incredible capabilities to those who have a spare 16-bit slot in their system. CPU, memory, and FPU are all on the card, and the host processor (286, 386, or i486) remains on the motherboard where it belongs, unaware of the coprocessor until it's told of it.

My review system, the 8140-PM, arrived with a 25-MHz 88000 CPU, instruction and data caches, and 16 MB of 32-bit memory. A 20-MHz 386 acted as the I/O server. A 330-MB ESDI hard disk drive and controller, a 150-MB cartridge tape drive, and an Artist TI12 graphics coprocessor card rounded out

the system. As the review unit included nearly all the options offered by Opus, I could exploit the system to the fullest.

The coprocessor card, the hardware star of the show, is an engineering marvel. The combination of main board and daughtercard can accommodate up to 24 MB of memory. As shown in the photo on page 170, the 88000 RISC processor is not one chip, but three in this configuration. The 88100 chip is the CPU, containing integer and floating-point circuitry, while a pair of 88200 memory management unit chips handles memory addressing and the dual caches. Power is drawn from the 16-bit connector; it is *continued*

Company
Opus Systems
20863 Stevens Creek Blvd.
Building 400
Cupertino, CA 95014
(408) 446-2110

Components
Processors: 25-MHz 88000; 20-MHz 386
Memory: 16 MB for the 88000; 4 MB
for the 386
Mass storage: 330-MB ESDI hard disk
drive, 1.44-MB floppy disk drive, 150-MB
cartridge tape drive
Display: 1280- by 1024-pixel, 19-inch
analog color display; intelligent graphics
coprocessor
Keyboard: 101-key PC compatible
I/O interfaces: Two serial ports; one
parallel port; game port; video port;
Ethernet port; 12 ISA bus slots

Price
Base system: $18,995
System as reviewed: $25,190

Inquiry 871.

also through this channel that all data
flows to and from the I/O processor.

Getting It off the Ground

Configuring the PM is easy—you don't.
The system arrived fully configured,
with every option tested and working.
When a power glitch or similar problem
caused the system to fail to boot, I called
Opus's technical-support department.
An individual who clearly knew her way

around the system walked me through re-
installing and restoring my backup tape
and helped me determine that the prob-
lem was a transient one. Opus's technical
support impressed me greatly.

There's more going on than meets the
eye with Opus's Unix implementation.
The heart of the I/O server is Opmon,
Opus's DOS-based I/O monitor pro-
gram, which handles the communication
between the Intel host processor and the
88000 CPU. It also routes all I/O re-
quests to the appropriate devices. Typing
unix loads Opmon and instructs the PM
card to test itself. Opmon then uploads
the Unix operating-system kernel to the
board from a DOS file. If the upload is
successful, Unix takes control.

Once Unix starts running, the Intel
processor keeps chugging away, looping
through the Opmon program and waiting
for I/O requests from the 88000. A DOS
file that Opmon reads at start-up main-
tains the machine's hardware configura-
tion. Changing the Unix configuration
rarely requires relinking the Unix ker-
nel, because Opus has compiled in sup-
port for all Opmon-supported devices.

Smooth Operator

As precarious as it seems, all these layers
work together smoothly. Opus loaded the
8140-PM with a hard disk drive, a tape
drive, serial and parallel interfaces, a
mouse, an Ethernet adapter, and a graph-
ics coprocessor. Under a normal load,
the 386 juggled the I/O requests admira-
bly. But once I ran tests intended to strain
the I/O subsystem, degradation came
swiftly and severely. (See "The BYTE
Unix Benchmarks," March BYTE.)

The most serious bottleneck resulted

when the X Window System ran while
some I/O-intensive task worked in the
background. When I accessed the floppy
disk drive or tape drive, response time in
X Window System fell to a disturbing
low. The problem was DOS. Its single-
threaded nature does not suit it to I/O
serving. It's a feat that Opus's engineers
were able to pull as much performance
out of DOS as they did.

The worst problem with I/O was that at
times the system seemed to freeze com-
pletely. The tape drive is the worst cul-
prit, forcing the rule of thumb that back-
ups and restores be performed while in
single-user mode. It's simply impossible
to get anything else done while the tape is
spinning, and the floppy disk drive has a
similar impact on performance.
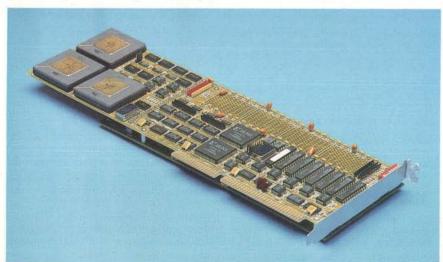
There can also be momentary pauses
in character output when large amounts
of cached data are written to disk, but
some simple tuning of the Opus Unix
kernel keeps this from getting in the way.
I look forward to Extended Industry
Standard Architecture– or Micro Chan-
nel architecture–bus implementations of
the Opus cards, and perhaps an OS/2 (or
even Unix) I/O server implementation,
so multiple I/O channels can be managed
more effectively.

The I/O layers are a great conve-
nience, allowing DOS and Unix to share
devices without hassle. The Unix system
can use a separate disk partition for its
data, but it can also build a file system in
a DOS file, so you can run both DOS and
Unix on the same disk without reparti-
tioning. Unix uses the host system's DOS
BIOS for all disk accesses, so virtually
any disk drive controller that works with
DOS will run with Unix, even if a DOS
device driver is required.

Trolling Through Unix

Opus's implementation of Unix, referred
to as Opus 5, is solidly built and per-
forms remarkably well. Those accus-
tomed (as I am) to running Unix on 386
systems will be pleasantly surprised by
the quality and speed of Opus 5. I'm
irked that Opus, like so many other com-
panies, chose to make Unix documenta-
tion an option, but any AT&T or Prentice-
Hall documentation covering System V
release 3.1 will suffice. Opus plans to
have release 3.2 available by the time this
review appears in print, and work on
System V release 4 is also under way.

In addition to the traditional System V
commands and utilities, Opus provides a
smattering of BSD Unix enhancements.
The C shell is there, of course, but the
real prize is the Berkeley socket library.
This set of functions handles interpro-

*The Opus 88000 coprocessor card. The three large, brass-centered chips are
the 88100 CPU and the dual 88200 cache/memory management units.*

Opus 8140-PM, 8110-PM

HIGH-LEVEL PERFORMANCE

Opus 8140-PM **13.5***

Opus 8110-PM **10.8**

Everex Step 386/33 **6.0**

	Opus 8140-PM		Opus 8110-PM		Everex
	Time	Index	Time	Index	Time
■ **C Compiler**	2.50	0.83	3.33	0.62	2.08
■ **DC Arithmetic**	0.25	2.52	0.30	2.10	0.63
▢ **Tower of Hanoi** (17-disk problem)	0.28	2.00	0.37	1.51	0.56
System Loading[1]					
1 concurrent background process	1.93	2.10	2.33	1.74	4.06
2 concurrent background processes	3.27	1.77	4.17	1.39	5.80
4 concurrent background processes	6.27	1.53	8.07	1.19	9.60
▢ 8 concurrent background processes	12.57	1.38	16.00	1.08	17.30

LOW-LEVEL PERFORMANCE

	Opus 8140-PM		Opus 8110-PM		Everex
	Time	Index	Time	Index	Time
▢ **Dhrystone 2** (without registers; Dhry./sec.)	50,385	3.64	41,218	2.98	13,487
Arithmetic (10,000 iterations)					
Arithmetic overhead	0.15	4.80	0.20	3.60	0.72
Register	1.95	1.50	2.48	1.18	2.92
Short	2.00	1.76	2.50	1.41	3.52
Integer	1.95	1.60	2.45	1.27	3.12
Long	1.93	1.62	2.48	1.26	3.12
■ Floating Point	3.83	3.11	4.80	2.48	11.92
Double	3.52	3.76	4.47	2.96	13.22
Throughput					
System call overhead (5 x 4000 calls)	1.98	0.56	2.57	0.43	1.10
Pipe throughput (read and write 2048- ×512-byte blocks)	1.40	0.66	1.80	0.51	0.92
Pipe-based context switching (2 × 500 switches)	0.80	0.79	1.02	0.62	0.63
Process creation (100 forks)	1.40	0.88	1.90	0.65	1.23
Excel throughput (100 execs)	2.07	1.66	2.67	1.28	3.43

13.5 **10.8**

Legend:
- ■ C Compiler
- ■ DC Arithmetic
- ■ Tower of Hanoi
- ■ System Loading
- ■ Dhrystone 2
- ■ Floating Point

Note: All results are in seconds, unless otherwise specified. Indexes show relative performance. For all indexes, an Everex Step 386/33 running Xenix 2.3.1 = 1.

[1] System loading was performed using Bourne shell scripts and Unix utilities.

*Cumulative index is formed by summing the indexed performance results for C Compiler, DC Arithmetic, Tower of Hanoi, System Loading (with eight concurrent background processes), Dhrystone 2, and Floating Point tests.

cess communications between processes not only on the host system, but also across network connections.

Once you start working with the PM, you quickly see that it's tuned for software development. The socket library could have been left out with no impact to users, but developers will find it a boon. Two other extras are of particular interest to this crowd: The DOS interface library, and the Diab C compiler.

The DOS interface library provides functions for transferring data to and from DOS media, as well as routines that

manage communications between Unix and DOS. Surprisingly, the library also includes a facility for executing DOS commands while Unix is still running. This interface exists on the command line, too. When you invoke it, the interface either executes a single DOS command and returns to Unix, or calls up a command shell and executes interactively until you enter exit.

Opmon is large, so any DOS program invoked via this method will have to fit into 256K bytes or so. If your host system has enough memory, it's possible to have

a full 640K bytes available through the use of IGC's VM/386 multitasking software. This works well enough, but most of Opus's disk I/O optimizations are lost when VM/386 is used. Normally, a DOS session suspends Unix, but there is an option to run DOS commands concurrently as well.

Typically, System V implementations use AT&T's Portable C Compiler. While it is competent for most applications development, it lacks the low-level, CPU-specific optimizations that can crank up

continued

performance-intensive programs to their peak. Opus responds to this by including an optimized RISC compiler from Diab that is a plug-in replacement for AT&T's. The difference is staggering: Dhrystone 2.0 numbers nearly doubled when optimized through the Diab compiler. Is this cheating? Hardly. Most compilers, including AT&T's, produce code intended for traditional complex-instruction-set-computer (CISC) CPUs. It's appropriate to have a compiler that is aware of pipelines and other RISC-specific traits.

Throwing the Workstation Switch

My review system included a graphics coprocessor card from Control Systems. To accompany it, Opus offers an X Window System port from Age, which uploads to the graphics card and runs native on the Texas Instruments 34010 processor. The Age software was in prerelease when I reviewed it, but it generally performed well. The zippy 88000 makes short work of the enormous overhead of X11 and OSF/Motif (a prerelease of which was also included), even at 1280 by 1024 pixels in 256 colors. I was able to open new windows—even the notoriously

sluggish xterm—more quickly than under any other X11 release 3 implementation I've seen.

The Age X11 port shows its power in drawing graphics, too. The X11.3's Achilles' heels, unfilled arcs and wide lines, popped onto the screen, and the Age port computes and displays full 256-color scanned images with rare speed.

I could have spent hours playing with xgif, a public domain image-display utility. Its best feature is the ability to adjust the size of an image when the window containing it is resized. This generally takes some time, but on the Opus system, most pictures simply snapped back into the resized window as soon as its new size was established.

Trail Blazer

The Opus PM series is a well-conceived, well-executed solution to the lack of affordable high-performance systems. The PM's ability to integrate into existing DOS-based systems extends the useful life of those machines, and Opus's preconfigured systems are both functional and reasonably priced. With growing interest in Unix, systems like these are

finding their way into places that, until recently, would have seemed unfit for such complex and capable computers. Opus has blazed a trail.

Opus subscribes to 88/Open, an organization that certifies that applications written to its specifications will run on any 88000-based machine. The 88000 is gaining ground, backed by an aggressive campaign by Motorola, and Opus provides a path by which even small developers can cash in on the future of RISC.

The cost of Opus's innovative design is I/O performance, but that doesn't diminish the value of the system. While the I/O bottleneck precludes the PM from acting as any but an informal file server, it is impressive as a stand-alone system.

For software developers and workstation users, this system is a bargain and must be considered seriously against anything the "big boys" have to offer in the same price range. Things may change, but for now, the Opus PM is a best buy for those who need big-iron performance. ∎

Tom Yager is a technical editor for BYTE. He can be reached on BIX as "tyager."

The Psion Mobile Computer
More than evolution. It's a revolution!

Designed to go. Anywhere. Any-time. Psion Mobile Computers are light (a mere 4.5 lbs), compact, and their slim, clam shell design is such that they invite being picked up and carried as a book.

Designed to be used. And used. And used. Models in the Psion Mobile Computer range have a typical battery life of 60 hours of continuous use. Couple this with an ergonomic design, a professional quality keyboard, and a high contrast, high definition graphic screen and you have a machine that's a pleasure to use. Anytime. All the time.

Designed to meet standards. And set them. Two ranges of Psion Mobile Computers are available. One runs industry standard PC Software. The other sets tomorrow's standards with a new graphics environment, a new built-in touch-pad to replace your mouse, voice inputs and outputs, and multi-tasking facilities.

Bristling with new technology— The Psion Mobile Computers have not been developed by squeezing more from conventional hardware. They are the result of new concepts, new production techniques and new technologies for memory storage, display, and voice processing. And the entire range of Mobile Computers use the Microsoft MS-DOS file data structures.

For more information on these and other fine PSION Products, contact:

XEC Products
13630 58th Street North, Suite 103
Clearwater, FL 34620
Phone: (813) 531-1422
Fax: (813) 530-5975

'Psion' and the logo are registered trademarks, and 'Psion Mobile Computer' and 'Psion MC' are trademarks of Psion plc.

'Microsoft' and 'MS-DOS' are registered trademarks of Microsoft Corporation.

More Than Upside-Down Mice

Five new trackballs
vie for the attention of
die-hard mousers

Bradley Dyck Kliewer

A mouse that's lying still and on its back may be one of two things: a late rodent or, if it's connected to a computer, a trackball. But despite the popular description of a trackball as an upside-down mouse, substantial differences exist in the designs and uses of the two pointing devices.

New trackballs have recently hit the market with sleek designs that capitalize on those differences. I'll evaluate five of these mouse alternatives: the CH Products RollerMouse, Kensington Expert Mouse, Logitech TrackMan, Micro-Speed PC-Trac, and Mouse Systems PC Trackball.

Trackball Basics

Both mice and trackballs use a ball to move a cursor, but you move the entire mouse to indirectly move the ball, while with the trackball, you move only the ball. In theory, trackballs should be more comfortable to use because your arm doesn't need to move as you move the cursor. In addition, a mouse requires a clear area on the desk on which to roam; a trackball remains stationary. But most trackballs are much larger than mice, and the footprint of the larger trackballs roughly matches the clear area required by a mouse. Also, I slide my mouse out of the way when not in use; the trackballs' rubber feet kept them from easily sliding out of the operating area.

All the reviewed trackballs use opto-mechanical sensors. That is, the external

Five new trackballs (clockwise from lower left): the Mouse Systems PC Trackball, MicroSpeed PC-Trac, Logitech TrackMan, CH Products RollerMouse, and Kensington Expert Mouse.

ball rolls against two internal sensing rollers: one mounted to sense horizontal movement and one mounted for vertical movement. The rollers, in turn, connect to a slotted or striped wheel, which opens or breaks two paths of light between LEDs and detectors. The roller mechanism is one of the most important factors in overall trackball performance, because a bad design can lead to rough action or slippage between the ball and rollers. In the latter case, the cursor's movement may be halting or jumpy. My test group showed differences in roller size, materials, and placement; thus, some trackballs performed better than others where these areas were concerned.

Also important are the number and placement of buttons, which let you highlight text, select a menu item, and perform other tasks. In this regard, I noted some marked differences between trackballs and the mouse, and even among the individual trackballs. I found click-and-drag operations particularly awkward

with a trackball. I'm obviously not alone in this opinion—most of the trackballs included some button-locking method. A few trackballs have one or two extra buttons for activating the lock, a feature I prefer. Button placement is important and one of the key areas that vendors use to distinguish products, but it's not an overriding factor: The trackball I liked best overall (the Mouse Systems PC Trackball) had the button placement I liked least.

Trackball resolution, usually rated in dots per inch, expresses the ratio of ball size to roller size and the number of slots in the wheel, which determines the number of on/off cycles per inch of movement. The higher the dpi, the less movement that's required to move the cursor. With the exception of the 300-dpi Logitech TrackMan, all the trackballs offered about 100 dpi at the detection level. However, the 100-dpi trackballs convert the signals to 200 dpi for compatibility

continued

MORE THAN UPSIDE-DOWN MICE

	RollerMouse	Expert Mouse	TrackMan
Company	CH Products 970 Park Center Dr. Vista, CA 92083 (619) 598-2518	Kensington Microware, Ltd. 251 Park Ave. S New York, NY 10010 (800) 535-4242 (212) 475-5200	Logitech, Inc. 6505 Kaiser Dr. Fremont, CA 94555 (415) 795-8500
Footprint	5 × 5 inches	4½ × 5½ inches	5⅜ × 4¼ inches
Hardware Needed	IBM XT, AT, or PS/2; Apple II or Macintosh; one disk drive; serial port (9-pin or 25-pin)	IBM PS/2 or compatible with pointing device port; 3½- or 5¼-inch disk drive	IBM PC, XT, AT, PS/2 (with available 9-pin or 25-pin RS-232C serial port), or compatible
Software Needed	DOS 2.0 or higher	DOS 3.0 or higher	DOS 2.1 or higher
Price	IBM XT, AT, and PS/2 model: $129.95 Apple II and Macintosh model: $119.95 PC bus model: $149.95	$169.95	Serial version: $139 Bus version: $149
	Inquiry 862.	**Inquiry 863.**	**Inquiry 864.**

with standard mouse drivers.

A trackball's driver software may modify the movement signal before it reports the results to the application. These modifications are known as *sensitivity*, *acceleration*, and *ballistics*. Basically, the driver software watches the signals coming from the trackball, and as the speed of movement increases, the driver may exaggerate the effect. This allows the cursor to cover large portions of the screen in a single, short, high-speed burst from your fingers. When you slow down, the driver returns to dot-per-pixel operation, or it may even drop signals to allow more detailed work.

If you will be using a trackball extensively, look for one that lets you tailor the acceleration curves. The process can be rather involved, but in the end your daily work will be more comfortable.

Test Track
Trackball "feel" is difficult to quantify. Our physical traits affect our opinions. In my case, I'm right-handed and have larger-than-average hands. I'm also accustomed to the mouse; I've been using one for six years.
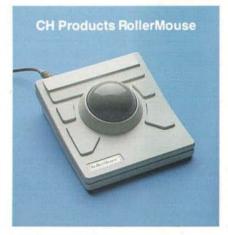
Four of my test trackballs used a serial port connection; the Kensington Expert Mouse used the PS/2 mouse port. I ran most of the serial trackball tests on an IBM AT with an Inboard 386, 6.5 megabytes of RAM, and a VGA display adapter. For the Kensington Expert Mouse and for the serial AutoCAD tests, I used

an IBM PS/2 Model 80-111 with 4 MB of memory and an 8514/A high-resolution graphics coprocessor. The AutoCAD tests ran on the 8514/A at 1024- by 768-pixel resolution. All other tests ran on the VGA display.

Software tests included AutoCAD release 10, Lotus 1-2-3 release 2.2, Lotus Symphony 1.2, Microsoft Word 5.0, OS/2 1.1 Presentation Manager, Microsoft Paint, and Microsoft Windows/386. I tried Aldus PageMaker 3.0 and Microsoft Excel under Windows. The trackballs passed all tests, with a few problems that I've noted in the individual reviews below.

I used Lotus 1-2-3 and Symphony to test the mouse shells, programs that allow trackballs to operate with character-mode-based applications that do not support mice. With the Windows applications, I tested general pointing, highlighting, and click-and-drag operations. The tutorial lessons in AutoCAD helped me develop a general feel for the trackballs. Microsoft Word was useful for button combinations, and it gave me a chance to try operations I've been using regularly for years. Paint gave me a feel for freehand drawing and cursor precision; the major test here was writing my name.

When I was not satisfied with a trackball's performance under a particular application, I adjusted the acceleration curve until I found the combination that felt best.

CH Products RollerMouse

- **Roller mechanism among the smoothest**
- **Best button design and placement**
- **Independent horizontal and vertical sensitivity adjustment**
- **Primitive Lotus 1-2-3 shell**

Under Paint, the RollerMouse created somewhat rough characters with a slight staircase effect. This was probably due to uneven 100- to 200-dpi translation rather than to ball slippage, which I did not notice. The RollerMouse felt smooth and silky and was quieter than the other trackballs. The 2⅕-inch-diameter ball was among the largest tested and, as a result, among the most comfortable to use.

I preferred the RollerMouse button placement and operation (four buttons in

PC-Trac

MicroSpeed, Inc.
44000 Old Warm Springs Blvd.
Fremont, CA 94538
(415) 490-1403

$3\frac{7}{8} \times 7$ inches

IBM PC, XT, AT, or PS/2 (with either serial,
bus, PS/2, or In Port interface)

DOS 2.0 or higher

Serial, PS/2, and In Port model: $119
Bus model: $139

Inquiry 865.

PC Trackball

Mouse Systems Corp.
47505 Seabridge Dr.
Fremont, CA 94538
(415) 656-1117

$3\frac{7}{8} \times 5\frac{1}{2}$ inches

IBM PC, XT, or AT (with available 9-pin or
25-pin RS-232C serial port)—Mouse Systems
plans PS/2 mouse connector version; EGA or
VGA card for Presentation Magician
(software included with the trackball)
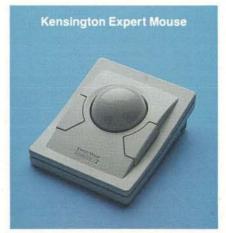
DOS 2.1 or higher

$119

Inquiry 866.

all, two on the left and two on the right) to those of the other trackballs, with the possible exception of the MicroSpeed PC-Trac's. The RollerMouse's upper buttons provide a click-lock, while the lower ones act as conventional buttons. The left and right button signals can be swapped via a switch setting on the bottom of the unit.

The RollerMouse produced garbage on the screen when I ran Word 5.0. Sometimes, the cursor masked screen characters or changed colors. The problem disappeared when I used the mouse driver included with Word 5.0.

With the RollerMouse, you can independently adjust horizontal and vertical sensitivity on a scale from 0 to 100. You can change the setting at any time from the DOS command line. You can also set the acceleration trigger level for double speed; however, acceleration is limited to normal or double speeds only.

The mouse shell program (with the Lotus 1-2-3 overlay) worked effectively with 1-2-3, but it required extensive modification to run with Symphony (no Symphony overlay was provided). The company did a good job documenting the menu-file format for those who want to make modifications. The 1-2-3 menu is one of the more primitive ones I've tested, limited to the menu key (/), the Edit key (F2), and the Enter, Backspace, and Escape keys. The button operation changes with the current 1-2-3 mode. Moving the ball moves the cursor within

the spreadsheet or menu.

The RollerMouse ball is hard to clean. Cover screws hide under the adhesive rubber feet (the manual suggests cleaning the surface of the ball without removing the cover).

Kensington Expert Mouse

- **Click-locking awkward under Word**
- **Metal rollers showed the most slippage**
- **Exceptional Lotus 1-2-3 and Symphony shells**
- **Acceleration adjustable from within an application**

The Expert Mouse sports a two-button design with a large ($2\frac{1}{8}$-inch) ball. You set right- or left-button operation via switches on the trackball's back or through the control-panel software.

You can plug the cable into either the left or the right side to accommodate either right- or left-handed users. The standard cable stretches a generous 8 feet.

The Expert Mouse does not have a click-lock button as such. However, you can set the trackball to activate click-lock by chording (pressing the buttons simultaneously), or you can even use the right or left button as a lock. I found this method annoying under Word, which uses all three of the button combinations. With chording set to lock the next button I pressed, Word required a double chord and then a third button press to activate the chord. But the click-lock was convenient with the other software. The Expert Mouse produced garbage under Word in graphics mode. As with the RollerMouse, using the Microsoft Mouse drivers solved the problem.

The Expert Mouse's trackball slipped more than any of the others. The Expert Mouse uses metal rollers (rather than plastic or rubber), and this may partially explain the slippage. Not surprisingly, I had trouble with fine control under Paint. Windows felt most comfortable at the lowest mouse sensitivity setting because of the rough handling.

The Expert Mouse includes pop-up control-panel software that lets you set button operation and acceleration parameters (on a scale of 0 to 9) while you're in an application. You can also adjust the acceleration curves with the MSPEED program. The curves from acceleration levels 1 to 9 can be individually tailored (the scaling factors affect the horizontal and vertical axes simultaneously).

The Lotus shells were among my favorites (they are nearly identical to the Mouse Systems menu shells). In 1-2-3 and Symphony, the trackball moves the cursor. The left button calls a mouse menu, the right button acts as the Enter key, and chording acts as the Escape key. The mouse menu has selections for activating the Lotus menu or choosing common commands (such as typing in mathematical operators and function keys).
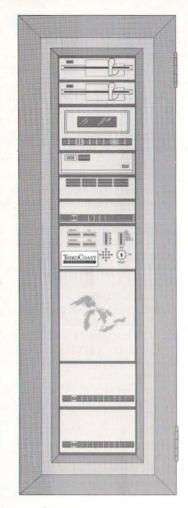
While these multiple keystrokes may seem awkward, I found the operation natural and flexible. I prefer the consistent button actions over shells that reduce the number of keystrokes with constantly changing button definitions. I would definitely use this shell under Lotus 1-2-3 and Symphony. In fact, the Kensington Expert Mouse's performance with these two products helps overcome the limitations of its two-button design.

Cleaning is easy: The ball simply lifts out of the casing.

continued

Logitech TrackMan

- **Unique thumb control of ball**
- **Excellent engineering—no ball skidding**
- **Favors right-handed users**
- **Lacks fine acceleration-curve control**

The TrackMan is unique in its use of the thumb as the 1⅖-inch ball's controlling digit (I controlled the others with my index and middle fingers). The index, middle, and ring fingers rest on the TrackMan's three buttons. Although Logitech doesn't include a click-lock, it was generally easy to hold these buttons for click-and-drag operations; still, my finger inadvertently released the button in some situations.

The design is decidedly for right-handed users, and as of press time, a left-handed model was not available. You can set the driver for left-handed use, but the instructions are fairly complicated, and even then, lefties must contend with buttons that sit on the left side and away from them. I found the thumb-based design the most difficult to control. After about an hour of practice, my accuracy improved, but I never did feel comfortable with vertical motions. Even with the trackball at the lowest sensitivity settings, it was hard to get fine control under Paint. Some people might prefer thumb control, but I am not one of them.

I didn't notice any ball skidding, however. In fact, Logitech's roller-mechanism engineering is excellent. I wish the designers had applied it to a more traditional layout. The rollers are soft rubber, and springs push them against the ball. The ball pops out easily with a push from behind for cleaning. (In terms of desktop real estate, the TrackMan requires a space approximately 5⅓ by 4¼ inches, making it one of the most compact units that I looked at.)

The TrackMan was the only reviewed model that did not follow Microsoft hardware conventions. You need Logitech-specific drivers for Windows and OS/2.

The control panel can pop up within an application to set sensitivity (from a scale of 0 to 10) and ballistics (off, low, or high). However, Logitech doesn't provide a way to fine-tune or customize acceleration curves. I might have liked the TrackMan better if I could have flattened the acceleration curves more at the lower levels.

The Symphony shell for the menu program is fairly well designed. However, I thought the 300-dpi TrackMan was much too sensitive for text-based work, even at its lowest settings. The product includes a Mouse-2-3 shell for Lotus 1-2-3 releases 1.x and 2.x (you must use the menu program for 3.0). I liked the shell: Common operations such as "sum all cells above" and "sum all cells left" are just a few keystrokes away, and the menus can be customized. But the shell remaps the gray plus and gray minus keys as right and down arrows to accelerate data entry. This does help data entry, but it can be annoying for formula entry.

MicroSpeed PC-Trac

- **Among the best for smooth operation**
- **Independent horizontal and vertical acceleration adjustment**
- **Convenient button placement**
- **Simplistic key-mapping software**

The PC-Trac I reviewed was a preproduction unit, although the circuit board appeared to be in final form and the documentation was typeset. [Editor's note: *BYTE normally restricts reviews to production versions of products. We made an exception in this case to provide a more complete product sampling and because our test unit lacked only cosmetic refinements.*] This is another 2⅕-inch ball, although the case provides for a smaller available surface than with the other trackballs of that size.

The PC-Trac's action ranks at the top, along with the RollerMouse, for smoothness. I especially liked the way this

design fit my hand. Also, the left and right buttons are within easy reach from any position, and they're contoured around the ball. However, with a footprint of almost 4 by 7 inches, the housing is large and may be too big for smaller-than-average hands. (My wife, who's also a computer user, tried PC-Trac, and she thinks that the housing is much too large for her hands.)

The click-lock is the middle of three buttons—a convenient position. You can also define the third button to act like the center button of a three-button mouse. The click-lock defaults to the left button, but it can be switched with a simple button-click sequence. The click-and-drag works fairly well even without the lock feature.

Fine control under Paint was a bit rough, but handling was smooth under the other applications. As with the RollerMouse, the Paint roughness may be due to the 100- to 200-dpi conversion. The driver software caused problems with the Word graphics screen, but once again, the Microsoft drivers solved the hang-up.

The PC-Trac's KEYMAP is a simple program that directly maps trackball movement and button clicks to specific keystrokes. You can select various mappings for several popular applications. However, when I loaded Symphony, KEYMAP became disabled.

I did not like the default selections for Lotus 1-2-3. The program's simplicity places severe restrictions on its abilities, in part because it doesn't offer the flexibility of pop-up menus or the ability to redefine button actions for different applications. At this point, the software makes the PC-Trac much better suited for applications already designed for a mouse.

You can tailor the acceleration curve to your own tastes. This is the only trackball that I tested that had complete and independent tailoring for the horizontal and vertical axes. The process is rather inconvenient, however. The SETBGT (set ballistic gain table) program modifies the device driver, so that you have to reboot the system to apply the new settings. You can select either your custom acceleration table or the default table from the DOS command line. You can also change the sensitivity from the command line (using a setting of 0 to 9 that selects a scaling factor for the acceleration curve).

The trackball cover is attached with three screws. Removing the cover to clean inside is a simple operation.

continued

Mouse Systems PC Trackball

- **Best overall performance**
- **No ball skidding**
- **Easily adjusted sensitivity controls**
- **Awkward button placement**

The PC Trackball uses a medium-size ball (1⅛-inch diameter), and only a relatively small percentage of the ball is exposed. Because its small exposed surface limits stroke size and, hence, movement, I found the default sensitivity setting too low. But I could easily adjust it by pressing Alt/Ctrl/left-mouse-button followed by a number (from 0 to 9) on the keyboard. The small exposed surface probably keeps skidding to a minimum (I didn't notice any) by directing most pressure downward onto the rollers.

Of all the trackballs I reviewed, the PC Trackball gave me the best cursor control and smoothest operation overall. I was able to produce a smoothly curved signature under Paint. I was not happy with its button placement and control, however. The left and right buttons are quite close to the ball. Clicking the center button was awkward. Furthermore, to use the center button as a click-lock, you have to hold the left or right button and then click the center button (reversing the order will not work). To release the click-lock, you have to click the center button again. Sometimes I forgot that the lock was on and thought the buttons had quit working. The company also offers left-handed drivers to let you swap button definitions.

Initially, I was not pleased with the response under AutoCAD. At the highest sensitivity level, it took too many strokes to move the cursor across the screen. But the acceleration curves can be tailored with the ULTRARES program. I found an acceptable setting by sharply accelerating the slower rates and leveling off acceleration at higher speeds.

The PC Trackball's menu and control software is nearly identical to the Kensington Expert Mouse's. The notable exception is 1-2-Click, which is a special Lotus 1-2-3 shell. 1-2-Click adds drop-down menus to 1-2-3 (although it does not contain all the 2.2 commands). The shell adds some useful features to 1-2-3 and can be disabled or reestablished with a click. I prefer the menu shell, but it's nice to have two choices, and I'm sure others will prefer 1-2-Click.

Despite the buttons, this was my favorite trackball. I think this comes down to the ball design, since I used the trackballs mostly to move the cursor around the screen and only occasionally clicked on the buttons. The ball and cursor movement was the most comfortable, and the acceleration was the best I tested: Stepping up from low to high speed was smooth. (The cover attaches with four screws, so the unit can be easily opened for cleaning.)

Final Tally

In the end, the Mouse Systems PC Trackball gets the highest marks for comfort, smooth operation, and cursor control. The MicroSpeed PC-Trac ranks a close second for comfort, and it may be my preference for CAD, Windows, or other applications that already provide support for a mouse.

The CH Products RollerMouse performed well in applications that didn't require a lot of detail or pixel-by-pixel control, including text-based applications. Personally, I wasn't a fan of the Logitech TrackMan's thumb-control design, but the 300-dpi unit worked well in my high-resolution, 8514/A environment.

I found the Kensington Expert Mouse to be the most limited of the five units, primarily because of its two-button design and the skidding ball. However, it worked fine under Windows, and it was especially good in Lotus 1-2-3 and Symphony.

None of the units in this new round of trackballs will lure me away from my mouse, which is proof that I'm a die-hard mouser. But if you're uncomfortable with a mouse, or if you're a beginning mouse user, these alternatives are worth a try. ∎

Bradley Dyck Kliewer is the author of EGA/VGA: A Programmer's Reference Guide *(New York: McGraw-Hill, 1988) and principal of DK Micro Consultants, a microcomputer consulting business in Bloomington, Indiana. He can be contacted on BIX as "bkliewer."*

HOW TO AUTOMATE A SMALL BUSINESS

by W. Gary Robertson

Automating a small business such as a doctor's office, accounting or legal firm can be challenging. Budgets often are limited and technical personnel non-existent. Having a reserve of computer hardware ready to support new employees is uncommon.

As these businesses grow, existing systems become strained. While larger organizations may be well served by a minicomputer, mainframe, or server-based network, these often are beyond the scope of a smaller business.

System cost, ease of use, training, and maintenance are important considerations. Multiuser systems, particularly DOS-based ones, typically perform best in each of these categories.

Multiuser systems save money by allowing one computer to support multiple users through terminals attached to the CPU. They also avoid the hardware expense and maintenance inherent in server-based LANs. DOS-based multiuser systems require minimum retraining, and allow employees to use familiar applications.

The automation of Dr. Susan LeGrand's medical practice illustrates how a multiuser system can affordably and easily computerize a small business.

When Dr. LeGrand established her practice she didn't own a computer. Paperwork quickly became impossible to manage, so she purchased an 80386 computer for insurance filing, accounting, patient records, and maintaining a large hospital census.

As her practice grew, Dr. LeGrand hired an assistant for her office manager. Dr. LeGrand considered purchasing a second computer and a LAN, or purchasing a multiuser operating system that would allow an inexpensive terminal to be a second workstation. The multiuser system cost $2,054 for the software, extra RAM and terminal, compared to $3,326 for the computer, interface cards and software for the LAN.

Dr. LeGrand chose The Software Link's DOS-compatible multiuser operating system, PC-MOS.

"Conceptually, the multiuser approach seemed ideal," Dr. LeGrand commented, "and when it was the least expensive, the decision was easy."

The system was installed over a weekend, avoiding office hour downtime. "Everything looked and worked the same," Dr. LeGrand said, "And we could continue to use our existing software and communication program."

System administration and maintenance is handled remotely by the PC-MOS distributor, J.S. Walker & Co. of Charlotte, NC.

"Having two workstations has really improved productivity," Dr LeGrand said. "And I can add up to three more workstations by simply installing RAM and terminals."

W. Gary Robertson is co-founder of The Software Link, Inc.

PC-MOS

The Multi-User Solution For The Multi-Dimensional Company

Odds are, you're part of a multi-faceted organization, one that's involved in many different projects and activities. Every day you juggle dozens of tasks. So why are your PCs still doing one thing at a time — for one person at a time?

Today's 286 and 386-based PCs provide the power to do much more. PC-MOS is the multi-user, multi-tasking software that unleashes that power, making your PCs as multi-dimensional as your business.

Minicomputer Power For The Cost Of A PC!

PC-MOS lets several users simultaneously run different programs on a single, high-performance PC. One user can run a spreadsheet, while another uses the word processor and several others access a database — all at the same time! So instead of replicating expensive PCs, each user has an inexpensive monitor or terminal. The benefits are lower cost, more control, better security and consistency across applications. And at $595 for a 5-user version, you can afford to get started today!

DOS Compatibility, NetWare Connectivity

PC-MOS lets users run the popular DOS programs they use now — even Microsoft® Windows 286. Our gateway to NetWare lets you expand your Novell network inexpensively and easily. And PC-MOS requires no expensive wiring, and no network management headaches.

Proven Reliable With 100,000+ Users

Because PC-MOS was the first DOS-compatible multi-user operating system, it offers broad compatibility and the reliability of time-tested software. More than 100,000 satisfied users trust their work to PC-MOS each day. Our latest version features an easy-to-use install program, lets you re-boot individual workstations, and supports high-resolution, bit-mapped color graphics.

Call us today. We'll show you how to add multiple dimensions to your PC.

THE SOFTWARE LINK

3577 Parkway Lane, Norcross, GA 30092
1-800-451-LINK, (404) 448-5465
FAX: (404) 263-6474   TELEX: 4996147 SWLINK

Circle 250 on Reader Service Card (DEALERS: 251)

VARS and RESELLERS:
Ask about our Sales Support Program

New CAD Test Shuffles 34010 Pack

Improved tests
and new drivers revise
the results of
our graphics board
roundup

Steve Apiki

Intelligent 34010-based graphics boards: The Artist Graphics Artist TI12 (left) and the Renaissance GRX Rendition II.
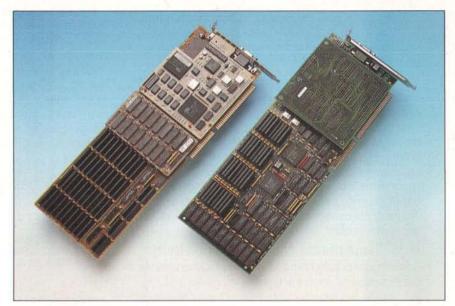
L ast November, the BYTE Lab evaluated AutoCAD and PageMaker performance of 34010-based intelligent graphics boards ("The Brains Behind the Graphics," November 1989 BYTE). Since then, we've improved our AutoCAD benchmarks and received new driver releases, so it's time to take a second look at the CAD results of these boards.

Three major changes have occurred since November. First, some of the manufacturers have released newer Autodesk Interface drivers, which allows us to test the boards more consistently. Our November tests used either ADI 3.1 or ADI 4.0 drivers, depending on availability. Our current tests use only ADI 4.0 drivers.

Second, we modified our test file. Our original drawing test had AutoCAD's OSNAP (object snap) feature turned on. Although this is not atypical, it added a large constant factor to each result and magnified the differences between ADI driver versions. Our new test drawing consists of last November's small AutoCAD drawing with the OSNAP parameter removed.

Finally, we relaxed the 512K-byte limitation on display-list memory size to simulate a more realistic high-end CAD environment. We tested each of the boards with as much on-board memory as possible and provided 4 megabytes of expanded memory for additional display-list storage.

The figure shows the results of these new AutoCAD release 10 tests. As in November, we split the times between boards with 1280- by 1024-pixel (high) and 1024- by 768-pixel (medium) resolutions. We tested all the boards in 256-color configurations with a Nanao FlexScan Model 9500 monitor. (Last November's top-performing high-resolution board, Number Nine's Pepper Pro1280, was unavailable for this review. Instead, we've included Number Nine's Pepper Pro1024 in the medium-resolution group.)

Among the high-resolution boards that we tested, the Artist Graphics (formerly Control Systems) Artist TI12 was a clear standout. The new TI12 GT driver, using on-board memory for display-list storage, pushed the Artist TI12 beyond the similarly configured Matrox PG-1281 and IMAgraph TI-1210-8. The Matrox board finished ahead of the IMAgraph unit, but the timings of all the boards were clustered fairly tightly. One oddity with the Artist TI12 was that, because of the display-list structure, you can't pan as far horizontally as you can on the other boards and still remain within the display list.

With its new software, Artist Graphics added more than top performance. The Artist TI12 has the sharpest configuration utility of the tested boards, and its control-panel interface to the display-list features is convenient. (For all the graphics boards, installation and use was straightforward, although the Pepper Pro1024's installation program is somewhat clumsy.)

The high-resolution boards have similar prices for similar configurations. All three boards generate a superior display and run trouble-free, but speed and a

1280- BY 1024-PIXEL RESOLUTION

Artist TI12
Artist Graphics
A Control Systems Company
2675 Patton Rd.
St. Paul, MN 55113
(612) 631-7800
As reviewed (includes 1.25 MB of video RAM, 3 MB of DRAM, and VGA module): $6685
Inquiry 854.

PG-1281
Matrox Electronic Systems, Ltd.
1055 St. Regis
Dorval, Quebec, Canada H9P 2T4
(514) 685-2630
As reviewed (includes 2 MB of Video RAM, 1.5 MB of DRAM, and Matrox PG-51-ACAD10/05 2.01 driver): $3945
Inquiry 856.

TI-1210-8
IMAgraph Corp.
11 Elizabeth Dr.
Chelmsford, MA 01824
(508) 256-4624
As reviewed (includes 1.25 MB of video RAM, 1 MB of DRAM, and IMAzoom 2.00 driver): $4585
Inquiry 855.

Note: Configurations are for 256-color operation at the maximum supported resolution.

1024- BY 768-PIXEL RESOLUTION

Advanced Graphics 1024 Board
Compaq Computer Corp.
20555 SH 249
Houston, TX 77070
(713) 370-0670
As reviewed (includes 1 MB of video RAM and 128K bytes of DRAM): $2098
Inquiry 857.

Cobra Plus HS
Vermont Microsystems, Inc.
11 Tigan St.
P.O. Box 236
Winooski, VT 05404
(802) 655-2860
As reviewed (includes 768K bytes of video RAM and 512K bytes of DRAM):
$3395
Inquiry 860.
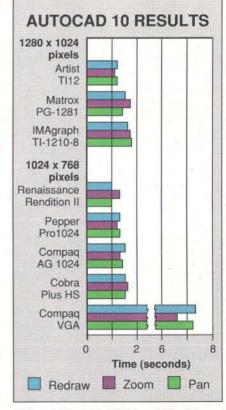
Pepper Pro1024
Number Nine Computer Corp.
725 Concord Ave.
Cambridge, MA 02138
(617) 492-0999
As reviewed (includes 1 MB of video RAM, 1 MB of DRAM, and Power9/ACAD 3.0 driver): $2940
Inquiry 858.

Rendition II
Renaissance GRX, Inc.
Cedar Park
2265 116th Ave. NE
Bellevue, WA 98004
(206) 454-8086
As reviewed (includes 1 MB of video RAM, 512K bytes of DRAM, and VGA module): $2695
Inquiry 859.

AUTOCAD 10 RESULTS

Although the Artist TI12 and Rendition II proved to be the top performers under AutoCAD release 10, the results put all these 34010-based boards in a relatively tight cluster. Benchmarks on the Compaq VGA, included here for comparison, illustrate the dramatic performance gain that a display-list processor can provide.

good interface make the Artist TI12 my favorite.
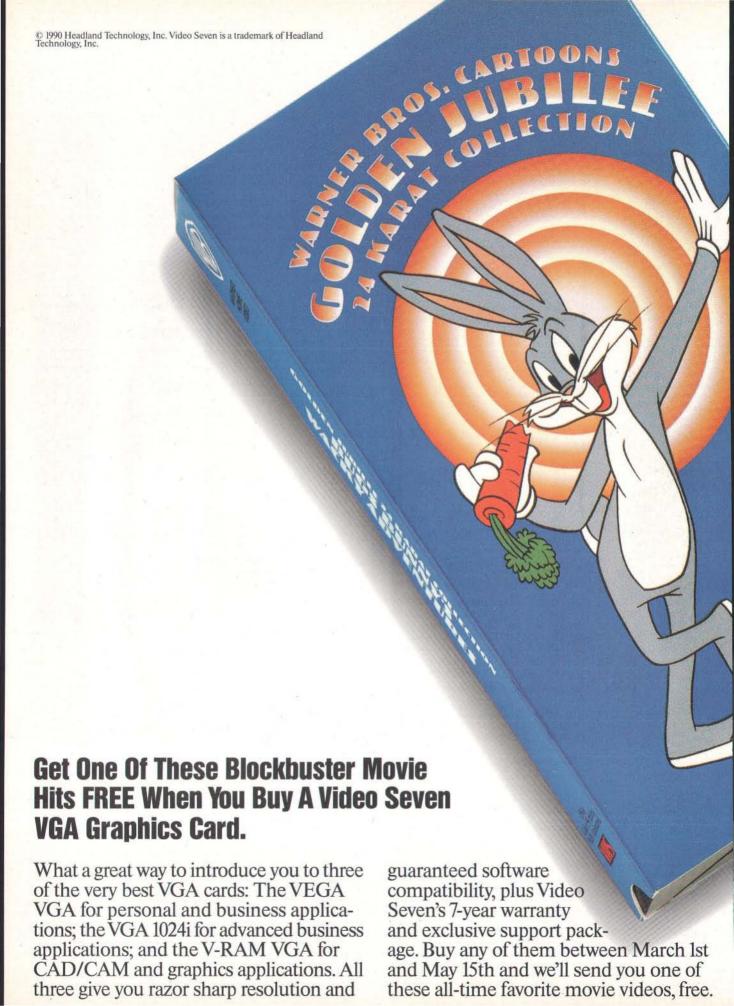
Renaissance GRX's Rendition II (see "Reviewer's Notebook," December 1989 BYTE) outpaced the rest of the medium-resolution pack. Here, too, the difference between the fastest and the slow-est boards was not great; all the boards finished within a half-second of one another.
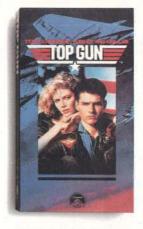
With a base price of $2095 (256-color version with 384K bytes of DRAM included), the Rendition II is the clear medium-resolution pick. All the boards turned in excellent performances, but the Rendition II offers it all at a very attractive price. ∎
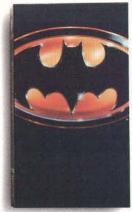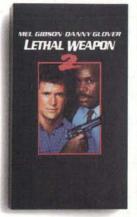
Steve Apiki is a testing editor for the BYTE Lab. You can contact him on BIX as "apiki."

Get One Of These Blockbuster Movie Hits FREE When You Buy A Video Seven VGA Graphics Card.

What a great way to introduce you to three of the very best VGA cards: The VEGA VGA for personal and business applications; the VGA 1024i for advanced business applications; and the V-RAM VGA for CAD/CAM and graphics applications. All three give you razor sharp resolution and guaranteed software compatibility, plus Video Seven's 7-year warranty and exclusive support package. Buy any of them between March 1st and May 15th and we'll send you one of these all-time favorite movie videos, free.

For every PC that has dreamed of becoming a CAD workstation.

© 1990 Hewlett-Packard Company   CPEG922

Macintosh CAD Comes of Age

Radius's display-list
processor helps Macs
match fast PCs for CAD

Jon Udell
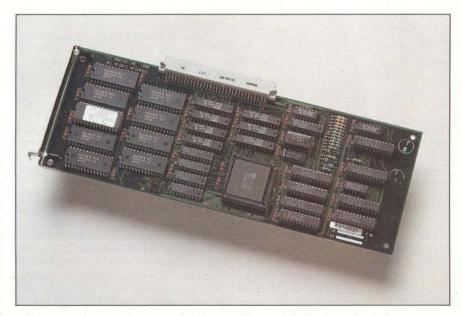
Blink and you've missed it. That's
how fast AutoCAD for the Macin-
tosh and the Radius QuickCAD
Graphics Engine can redraw the BYTE
Pantheon, our architectural modeling
test. As a result, Radius's $1495 NuBus
board answers CAD users' pressing need
for a Macintosh-based display-list pro-
cessor (DLP).

Ironically, CAD has been the Achil-
les' heel of the otherwise graphically
sophisticated Macintosh. A fast 386 PC
arguably makes a more cost-effective
CAD workstation than a 68020 or 68030
Mac. Equip the PC with a 34010-based
coprocessor, and it can blow the doors
off any Mac on pan, zoom, and redraw
operations. Now, there's an equalizer.
QuickCAD, with its VL86C010 integer
RISC chip, manipulates display lists
(sets of line segments or vectors) and ren-
ders them into the Mac's frame buffer.

Currently, QuickCAD works with
three Macintosh CAD programs: Auto-
desk's AutoCAD, Versacad's VersaCAD
Design, and IGC Technology's Pegasys.
QuickCAD's code runs in the 256K
bytes of static RAM on the board. The
display list resides in system RAM, and
it can overflow to disk, Radius says.

QuickCAD's benefits and limitations
match those of PC-based DLPs. A dis-
play list is a special kind of cache (see
figure 1). It stores endpoints that define
the line segments of a CAD image. Ab-
stract shapes rather than screen coor-
dinates represent those points. The pro-
cess of building that virtual display—a

regen—requires traversal of the CAD
database and, in the case of three-dimen-
sional CAD programs, a laborious map-
ping that projects a 3-D model onto a 2-D
display plane. A regen proceeds slowly
and, because it rewrites the display list,
gains nothing from a DLP. But once
you've built a display list, the DLP kicks
in. There is no need to revisit the data-
base or repeat the vector calculations.
Want to redraw the image? The DLP just
rattles through the list of cached vectors
and redisplays them. Pan or zoom? The
DLP clips or scales the vectors on their
way to the screen.

Life's great until you invalidate the
cache by zooming beyond the limits of
the virtual display or rotating a 3-D
model. Since most popular CAD pro-
grams handle 3-D models, that's an im-
portant restriction. A DLP won't enable
your PC or Mac to spin fully rendered
teapots, Volkswagens, or Pantheons in
real time. That can't happen until the
CAD programs maintain 3-D display
lists that FPU-assisted graphics copro-
cessors can manipulate. DLP vendors

don't emphasize this restriction, for ob-
vious reasons, and it's not as well under-
stood as it should be. So caveat emptor:
A DLP like QuickCAD can come in
handy, depending on the kind of work
you do, but it's not a plug-in 3-D engine.

Tips and Tricks
I tested the QuickCAD on a 5-megabyte
Mac II equipped with two monitors. The
primary monitor was a 19-inch Radius
Color Display driven by a DirectColor/
24, Radius' 32-Bit QuickDraw board.
The second monitor was a standard 13-
inch Mac II display with an Apple 8-bit
video board. QuickCAD comes with an
INIT/cdev, which loads the RISC pro-
cessor's static RAM with the code that
traverses a display list, transforms vec-
tors, handles NuBus arbitration, and
shoves pixels into a frame buffer (or, in
the case of multiple monitors, several
frame buffers). The cdev provides two
controls—a master switch and a Caps
Lock switch. When you enable the latter,
you can toggle QuickCAD on and off
continued

interactively with Caps Lock. It's a handy way to gauge the effect of the board.

QuickCAD is a superset of Radius's QuickColor accelerator board (for more information about the accelerator, see the text box "The QuickColor Option" on page 194 in the December 1989 BYTE). Because QuickCAD subsumes the QuickColor accelerator, there's also a QuickColor INIT/cdev. QuickColor's INIT works like QuickCAD's, except that its job is to replace certain Quick-Draw bottleneck routines. Applications that rely exclusively on QuickDraw needn't know about QuickColor. Its

benefits—fast scrolling of deep-color and bit-mapped images—come automatically. QuickCAD, by contrast, operates outside QuickDraw's scope. A program wanting to export a display list must know a DLP is available to handle it.

As it turned out, a QuickColor board already resided in my Mac II, and its IN-IT/cdev still sat in the System Folder. So QuickCAD's installer overwrote it and added the QuickCAD INIT/cdev. I re-booted the Mac, turned QuickCAD on, fired up AutoCAD, loaded the Pantheon model, and issued a redraw command. No dice—it wasn't any faster than a nor-mal redraw.

After a phone call to a beta user, I tried again without the QuickColor board (it fights with QuickCAD and is redun-dant, anyway). Still no dice. Spelunking around, I found that the primary moni-tor, the Radius Color Display, was set for 16-bit pixel depth. I dropped back to 8-bit pixel depth (as, in fairness, a READ-.ME file instructs you to do) and tried once more. That did the trick: Redraws, pans, and zooms became nearly instan-taneous (see the photo on page 190).
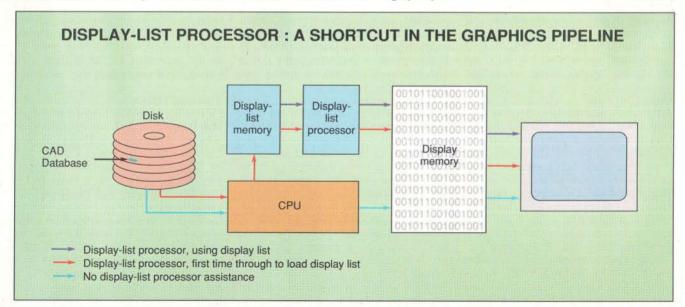
Figure 2 tells an interesting story. Running BYTE's AutoCAD benchmark, an unassisted Mac II can't hold a candle to a VGA-equipped Compaq 386/20. When you equip both machines with DLPs (I used the 34010-based Advanced Graphics 1024 on the Compaq), the Mac runs the redraw, pan, and zoom tests nearly as fast as the PC. QuickCAD ac-celerates VersaCAD and Pegasys by a

comparable factor, according to the ven-dors of those programs. Note, however, that the Compaq retains more than a fourfold advantage over the Mac on the regen test, with or without a DLP.

One thing puzzled me. On the PC, AutoCAD's video-configuration menu includes an AutoCAD Display Interface option that directs the program to use a DLP. That option doesn't appear in the Mac version of AutoCAD, and Quick-CAD didn't come with any application-specific drivers. How did AutoCAD link up with the Radius accelerator? The answer suprised me: QuickCAD's in-staller had silently patched AutoCAD's INIT. Radius isn't happy with that scheme, according to a product manager I spoke with. The other two supported programs, VersaCAD Design and Pe-gasys, will come bundled with Quick-CAD INITs, and Radius hopes that AutoCAD will do the same.

The personalities of the two Radius boards—QuickColor for bit maps and QuickCAD for vectors—aren't really complementary. A CAD program that does true-color rendering (Intergraph's MicroStation, which doesn't yet support QuickCAD, is one such program) isn't bound by bit-transfer rates, as are image-processing programs. Sorting polygons and computing light intensities take much longer than pushing pixels—even very deep ones—to the screen. Quick-CAD does not include QuickColor for the sake of CAD programs. It's there
continued

DISPLAY-LIST PROCESSOR : A SHORTCUT IN THE GRAPHICS PIPELINE

→ Display-list processor, using display list
→ Display-list processor, first time through to load display list
→ No display-list processor assistance

Figure 1: *A display list is a special kind of cache that stores endpoints to define the line segments of a CAD image. Once you've built a display list, there's no need to revisit the database or repeat vector calculations. This makes most pans, zooms, and redraws almost instantaneous.*

because extending QuickColor's code, essentially an operating system for the RISC board that can multitask with the Mac operating system, was safer than replacing it.

The net result is that QuickCAD lets you do everything that a QuickColor board does. I loaded a 24-bit-deep color photo into PageMaker and observed noticeably faster scrolling with QuickColor enabled. Architects and others who create intricate CAD models and glossy

Figure 2: *BYTE's AutoCAD benchmark tells the story: When both machines are unassisted, the Compaq 386 wallops a Mac II. But the Macintosh/QuickCAD duo stands shoulder to shoulder with a Compaq 386/34010 combo in those operations that can benefit from a display-list processor.*

When running QuickCAD under AutoCAD, redraws, pans, and zooms on the BYTE Pantheon model became nearly instantaneous.

presentations will find both halves of the board useful. Note, however, that you have to switch between 8-bit and 16-/24-bit color modes.

It's conceivable, of course, that you'll want fast bit maps and fast CAD at the same time. Imagine a compound document that's a mixture of richly formatted text and rendered CAD images. Applications don't do that yet, but they probably will, so Radius plans to make QuickCAD compatible with at least 16-bit color. What if you already own QuickColor but now find that you require a CAD accelerator? Radius says that it's considering an upgrade for the difference in price between the two.

Newton's Apple

Initially, QuickCAD enjoys only limited software support. Will more Macintosh CAD vendors embrace the board? Undoubtedly, although it'll be harder for some than for others. Programs that already maintain display lists and don't depend heavily on Mac-specific optimizations will convert easily. (Significantly, none of the programs that now support QuickCAD grew up on the Mac.) Others face a more difficult challenge. Intergraph's MicroStation falls into this category. It images directly into off-screen buffers, which makes the program responsive but leaves no convenient hook for a coprocessor. Intergraph and Radius are studying how to proceed.

The biggest obstacle, of course, is the lack of 2-D and 3-D graphics standards. Thus far, Apple has protected Mac users from the hair-raising incompatibilities that plague the PC world. For example, Mac programs written for 8-bit color generally work transparently with 32-Bit QuickDraw. Mac users take this for granted, which amazes PC folk. Apple's dictatorship has clearly paid rich dividends. But the Mac hasn't done well in the engineering market. Hardware innovators like Radius now hold the key to Apple's future there.

QuickCAD, and products like it, will surely help sell Macs as CAD workstations. That'll mean a certain loss of innocence, though, as application/driver/board combinations multiply. Apple should embrace QuickCAD and, as with 32-Bit QuickDraw, move quickly to establish a standard that fosters creative diversity among hardware vendors while preserving the simplicity that makes the Mac great. ∎

Jon Udell is a BYTE senior technical editor at large. He can be reached on BIX as "judell."

Help for the C Sick

Microsoft BASIC PDS
7.0 may be the antidote
you're looking for

Mike Wiggins

Microsoft has once again given us what seems to be a natural evolution from its QuickBASIC environment. The Microsoft BASIC Professional Development System 7.0 is packed full of new and enhanced features to whet every BASIC programmer's appetite. BASIC PDS 7.0 enhancements include the Microsoft Quick-BASIC Extended (QBX) environment, support for OS/2, EMS 4.0 support, add-on financial libraries, user interface toolboxes with source code, and the new Microsoft Professional ISAM (indexed sequential-access method) database engine. Compiler junkies will like the improved code generation, which creates tighter, faster code. Heck, even C snobs will find this package interesting.

Microsoft has finally listened to professional BASIC programmers who've been demanding professional-quality development tools for years. Microsoft endowed PDS with a feature-rich environment for fast code development without loss of program performance. Many of the limiting factors typically associated with BASIC have been addressed in this package. For example, you can now have larger code and data sizes, do string moves into far memory, and make .EXE files as large as 16 megabytes. The new language enhancements provide greater control over DOS files, static arrays, error handling, and the communications port. BASIC PDS 7.0 also comes with the Microsoft Professional Editor and the

Microsoft BASIC PDS 7.0 adds, among other things, a mouse-driven, character-based user interface.

CodeView window-oriented debugger. If you are serious about developing big BASIC programs but have been running into limitation problems with other compilers, then you should take a look at BASIC PDS 7.0, because it can make your code jump through hoops.

The main thrust of the package is toward the business application programmer. BASIC has evolved from a beginner's language into a more structured and modular one, making it ideally suited for group programming projects. Companies with programming staffs will find BASIC PDS 7.0 well suited for many of their programming requirements. Although I wouldn't throw out my C compiler just yet, or translate all my C code to BASIC, I would seriously consider writing some of my small- to medi-

um-size projects in BASIC using this compiler. While BASIC PDS 7.0 has an excellent help system, it does not come with a tutorial, so the novice and hobbyist would be better off sticking with QuickBASIC 4.5 for experimentation.

QBX Environment
For those not familiar with the Microsoft QBX environment, it's what Microsoft calls an "instant development environment," which uses a threaded p-code technology. The environment offers the developer integrated editing and debugging for fast code development. QBX has a pull-down menu interface that can be controlled via the keyboard or an optional mouse. This is a full-screen editor with all the editing features needed for

continued

BASIC PDS 7.0 VS. QUICKBASIC 4.5

*BASIC PDS 7.0's I/O performance is improved over previous compilers',
particularly with commonly used sequential files. (All times are in seconds.)*

Sequential disk I/O on a 1-MB file

Record length (in bytes)	BASIC PDS 7.0		QuickBASIC 4.5	
	Write	Read	Write	Read
1	689.26	1084.12	787.36	1178.92
10	103.47	136.66	136.21	169.50
100	38.62	39.05	71.40	71.40
1000	34.77	38.50	70.97	70.85

Random disk I/O on a 10K-byte file

1	32.84	32.19	33.94	33.34
10	3.36	3.24	3.40	3.41
100	0.49	0.39	0.39	0.38
1000	0.16	0.22	0.16	0.17

fast code production. Once you've completed your program, you select the Run option on the menu to compile it in memory or into an .EXE disk file. An integrated debugger makes it easy to correct any compilation or program errors. The environment provides you with greater capability for system development. By shrinking the environment down with a "No Frills" switch, you can create megaprograms of up to 16 MB while inside the QBX environment.

Don't let the QBX environment fool you. This is a serious compiler, and it can easily handle large database applications like customer tracking, inventory control, and publication mail listings. Microsoft set out to produce an environment for serious developers, and it has created a monster language to rival even the great C serpent for speed and power.

BASIC PDS 7.0 has a complete on-line documentation system managed by the Microsoft Advisor help system. This context-sensitive help system has hypertext links that let you move quickly among key concepts. This is truly a complete system, and if you can find it in the manuals, you will find it on-line. You can even add your own help files to the system with the HELPMAKE utility.

Simply BASIC
The compiler has been improved in many areas: math, I/O, event handling, code generation, and optimization. Microsoft claims that these enhancements can increase the speed of your programs to up to 1900 percent and can decrease code size to up to 68 percent over QuickBASIC

4.5. Although the code I tested did not have anything near that kind of improvement, it did run faster, and the .EXE file sizes were noticeably smaller.

The math portions of BASIC have been beefed up in a few spots. If the math package senses the coprocessor, it will use it. But even without the hardware assist, the package is quite capable of handling floating-point math with 64-bit precision. All the rudimentary math programs I threw at it ran very fast. If you develop only for machines with math coprocessors, you can compile with the coprocessor-only option to eliminate the emulation package, thus reducing the size of the .EXE file.

Getting data in and out of any program can often cause bottleneck problems. BASIC PDS 7.0's I/O routines have been fine-tuned as part of the upgrade. I tested both sequential and random-disk-access writing and reading from a 1-MB file with record lengths of 1 to 1000 bytes. BASIC PDS 7.0 access is faster than QuickBASIC 4.5 during sequential access, but it was only marginally better during random access (see the table).

There is built-in screen support for a variety of adapters, including the Olivetti color adapter, Hercules graphics card, MCGA, CGA, EGA, multicolor graphics array, and VGA. It supports the VGA and MCGA adapters in 320- by 200-pixel resolution with 256,000 colors in mode 13. Again, there is a slight improvement in output to the screen over QuickBASIC 4.5, but nothing mind-boggling.

The event handling no longer has to check for events after executing each

line. By using Event On and Event Off, you can minimize the amount of checking done in your program.

Microsoft has improved the code-generation algorithms called "Thunking" and "Peep-holing." Thunking looks for frequently occurring code patterns, and then it generates an assembly subroutine in near memory for quick access. Peepholing looks for local string-related code patterns for optimization. These routines contribute to the compiler's overall ability to generate executable files that are small and fast.

A few new command-line switches have been added to the BASIC compiler's stand-alone operation. The /G2 switch tells the compiler to generate optimized code for the 286 processor. This lets your program take advantage of the 286's expanded instruction set. The /OT switch is used to optimize the calling procedure for your SUB, FUNCTION, and DEF FN routines. The /Fs switch is for far string operations, and only objects compiled with this switch set are compatible. The /Ix: switch controls memory management for the new ISAM database.

On the Financial Front
Some of BASIC PDS 7.0's new features offer greater power specifically to the business application developer. The new library routines come in both standard (.LIB) and Quick (.QLB) library formats. These financial libraries consist of formulas, numeric formats, and serial number generators. The financial library (FINANCER.QLB) offers commonly used business formulas, such as depreciation, interest, rate of return, and principal and interest payments. The date/time library (DTFMTER.QLB) delivers a way to generate serial numbers from date and time arguments. Finally, the format library (DTFMTER.QLB) returns properly formatted strings from numeric data types, including the new Currency data type.

The Currency data type is perfect for accounting and other financial applications. The currency number is stored as an 8-byte integer, and it can represent numbers with a range of $\pm 9.22E14$ with 19-digit accuracy. Now you can represent money the way you're used to seeing it: in dollars and cents.

Application Toolbox
The Matrix Math, Presentation Graphics, and User Interface toolboxes can give your programs some pizzazz. These toolboxes come with source code that you can use in your program development.

continued

However, Microsoft notes that "these are examples only and should not be relied upon as a fully tested add-on library." You might want to give the programs the once-over before you use them.

The Matrix Math toolbox demonstrates how math can be performed on two-dimensional matrices. Gaussian elimination, determinant and inverse of a square, and the BASIC math functions are included in the toolbox. You can perform these operations on most of the numeric data types in BASIC PDS 7.0.

The Presentation Graphics toolbox produces a variety of business-style graphics. Whether you need a scatter chart or a simple pie chart, there is sure to be a toolbox routine you can draw on. What good are charts without words? The Font toolbox is included as part of the Presentation Graphics toolbox. Two fonts are supplied (Helvetica and Times Roman), and you can also use Windows/286 bit-map .FON files with the Font toolbox.

The User Interface toolbox should be very popular. It gives your application a professional-looking user interface with character-based windows, pull-down menus, and mouse control. The interface has many of the features of the QBX environment. All three of these toolboxes will save the professional business programmer a bundle of development time.

The ISAM on the Cake

The Microsoft Professional ISAM database engine is simply the most amazing addition to the BASIC language that I have seen in years. It seems to be a com-

plete implementation of ISAM. The last time I worked with ISAM was in COBOL a few years back. This is a very good implementation of ISAM, and professional programmers who need database functionality in their programs could save months of development time by using it. Mind you, it has a bit of overhead, so it is not the answer for all your database problems.

The smallest ISAM data file generated is 64K bytes, but it has the power to handle larger database files of up to 128 MB in size. That 64K-byte initial file size has 39K bytes reserved for file management, leaving 25K bytes for new indexes, tables, and data. When the file needs to expand, it allocates 32K bytes at a time.

All the information for the handling of the database is stored in one file along with your data. The ISAM file is a series of tables and indexes. The tables consist of the actual data records, and the indexes provide quick access to table data. You can create an index on any field or combinations of fields in a table. Multiple tables can be stored in one database file offering greater flexibility for applications. Records are stored in a table in the order in which they were entered. If you want to access a record differently, simply create an index for that table that changes the access order or record key.

There are a few ways that you can install ISAM for use with your programs. Microsoft supplies two TSR-style programs, PROISAM and PROISAMD, or you can link your program directly with the ISAM library. PROISAM has full access to the database but lacks the ability to create or delete indexes and tables. PROISAMD has no file access restriction. By running a TSR, you can have access to the ISAM engines while in the QBX environment. This lets you develop and test your code without having to link in the large database library.

Some new keywords added to the BASIC language support ISAM. These keywords are a smooth and natural extension to the language. Take the OPEN statement, for example: `OPEN database$ FOR ISAM tabletype tablename$ AS filenumber%`. This is a natural replacement for a regular OPEN statement. Microsoft provides useful ISAM program examples in source code.

BASIC PDS 7.0 has four ISAM utilities for importing and maintaining ISAM files. One of the biggest problems with creating a large database is ensuring its integrity. With all that complex data crammed into a single file, the database structure is quite fragile. Two of the utilities, ISAMREPR and ISAMPACK, give

you a way to repair corrupted database files and optimize the size of the database file for better performance. The other two utilities allow you to convert your old database files. These programs can convert dBASE, Btrieve, ISAM, and ASCII files. If you're looking for a little more power from your database, you might want to consider BASIC PDS 7.0.

Pithy Stuff

Overall, I like the new QBX environment. The pull-down menus provide great control over the environment. Although you don't need a mouse, it comes in handy when using the menus or moving chunks of text around in large programs. The fast response of the QBX interpreter let me create and debug my test programs quickly. The Microsoft Professional ISAM engine is useful for creating any business applications that require fast access to stored data.

I had no trouble running BASIC PDS 7.0 on an 8-MHz 286 AT compatible with 1 MB of memory. If you're planning to write ISAM applications, I recommend a faster 286 or 386 with a minimum of 2 MB of EMS 4.0 memory and a 40-MB hard disk drive. QBX and the ISAM engine are geared to use EMS 4.0 memory, if it's present. BASIC PDS 7.0 also generates OS/2-compatible code, but because of my system's configuration, I was unable to test this capability. Some of BASIC PDS 7.0's operations—most notably the ISAM engine—are not supported under OS/2, and there's no direct support for Presentation Manager.

Satisfy Your BASIC Urges

Microsoft BASIC PDS 7.0 is right on target in my book. Microsoft has produced another winning version of BASIC, and it's about time. BASIC PDS 7.0 is a feature-rich package with all the speed and versatility to rival any other high-level language.

If you haven't used BASIC in a while, take a look at it now. You'll be pleasantly surprised. I would recommend this package to anybody who needs to develop clean, structured code quickly. Both business and scientific application programmers will benefit from all the improvements in BASIC PDS 7.0. So, if you're tired of writing dBASE code, or just plain "C sick," reconsider working in BASIC. ∎

Mike Wiggins was a founding member of Multimate and currently operates Hindsight, a consulting firm in South Windsor, Connecticut. He can be reached on BIX as "mwiggins."

Find Out Why We're Rated No.1 for Service & Support.

"PC Brand is the LL Bean of personal computer mail order...
...PC Brand wants no unhappy customers,...it's service and support policies help to insure that."

286 SYSTEMS FROM $599

PCBRAND 286/12 _____ $599

12 MHz Clock, Zero Wait Operation,
Norton SI 15.3 Landmark ™ Speed 15.1MHz,
512K RAM, 1.2MB or 1.44MB Drive, 101-Keyboard,
2 Serial and 1 Parallel Ports

PCBRAND 286/16 _____ $749

16 MHz Clock, Zero Wait Operation,
Norton SI 19.0 Landmark ™ Speed 20.6MHz,
512K RAM, 1.2MB or 1.44MB Drive, 101-Keyboard,
2 Serial and 1 Parallel Ports

PCBRAND 286/20 _____ $899

20 MHz Clock, Zero Wait Operation,
Norton SI 23.0 Landmark™ Speed 26.7MHz,
512K RAM, 1.2MB or 1.44MB Drive, 101-Keyboard,
2 Serial and 1 Parallel Ports

Standard System Features:

- 80286-12, 80286-16, 80286-20 operating at 12 MHz, 16MHz, or 20MH z w/Zero Wait
- 512K RAM expandable to 8MB on the System board using 256K or 1MB 100ns RAM
- 1.2MB 5.25" or 1.44MB 3.5" Diskette Drive
- FCC Class "A", Intended for business use
- High performance 16bit VGA Cards with 1024x768 capability on all VGA systems
- 1:1 Interleaved Hard/Floppy Drive Controller, 1Mb/Second disk transfer rates on all 100Mb drives or larger
- Enhanced 101-key Click/Tactile Keyboard
- 2 Serial & 1 Parallel ports on std-configurations
- High Capacity System Power supply
- Real Time Clock/Calendar with 5 Year Battery
- 80287 Co-Processor Support
- AMI BIOS w/full MS/DOS, OS/2, XENIX, UNIX, NOVELL, 3COM and PCNET compatibility
- Built-in System Board LIM 4.0EMS hardware
- User configurable I/O timing permitting compatible operation w/older peripherals or faster I/O for newer devices
- 8 Slot motherboard design (5 16Bit & 3 8Bit)
- Medium foot print case w/6 Disk Drive bays

Options:

- Low profile Slim Line Case w/3 Disk Bays available at no extra charge (pictured above)
- Mini Size desk top Tower ® Case w/4 bays
- LCD or Plasma Portable
- Factory Installed RAM Upgrades
- Custom configurations w/Name Brand peripherals of your choice

PC BRAND 286/12

w/512k, Hard Disk Drive, Monitor & Video Card

Hard Drives: Mb/Ms	20/40	40/25	66/25	100/25
No Video	$929	$1029	$1209	$1349
Mono	$1049	$1149	$1329	$1469
VGA-Mono	$1229	$1329	$1509	$1649
VGA-Color	$1459	$1559	$1739	$1879
SVGA/Color	$1569	$1669	$1849	$1989

PC BRAND 286/16

w/512k, Hard Disk Drive, Monitor & Video Card

Hard Drives: Mb/Ms	20/40	40/25	66/25	100/25
No Video	$1079	$1179	$1359	$1499
Mono	$1199	$1299	$1479	$1619
VGA-Mono	$1379	$1479	$1659	$1799
VGA-Color	$1609	$1709	$1889	$2029
SVGA/Color	$1719	$1819	$1999	$2139

PC BRAND 286/20

w/512k, Hard Disk Drive, Monitor & Video Card

Hard Drives: Mb/Ms	20/40	40/25	66/25	100/25
No Video	$1229	$1329	$1509	$1649
Mono	$1349	$1449	$1629	$1769
VGA-Mono	$1529	$1629	$1809	$1949
VGA-Color	$1759	$1859	$2039	$2179
SVGA/Color	$1869	$1969	$2149	$2289

®Tower is a registered trademark of NCR Corporation

Desktop Supercomputing

About 20 years ago, in the research labs at IBM, computer scientists were working on a 1-foot cube that in 10 years, they said, would contain the computing power of a room full of computers. We were the latest crop of new hires, and we didn't think it was possible. We were wrong. Nearly 10 years ago, a descendant of that box became the IBM PC.

Today, the desktop supercomputer is becoming a reality. Bit by bit, microcomputers are chipping away at the power of the supercomputer by borrowing its techniques and concepts. In "Separated at Birth," Bob Ryan discusses areas in which microcomputers are approaching the big boys, and areas they have not yet conquered.

One of the techniques used extensively in supercomputers that holds great promise for microcomputers is multiprocessing. You've seen the announcements for the Compaq Systempro, the Zenith Z-1000, and others that contain more than one CPU. But you don't necessarily have to go out and buy a multiprocessor system to get multiprocessing power. If you're on a network, you already have access to a supercomputer. In "Spending Your Free Time," David Gelernter and James Philbin explain how to use the idle machine cycles of a network's many computers for parallel processing.

Parallel processing is often used to describe any kind of multiprocessing. But true parallel processing is a refinement of it—a subset, perhaps. Multiprocessing refers to the simultaneous processing by more than one processor of different applications, programs, or processes within a program. True parallel processing is the simultaneous execution by more than one processor of different instructions within the same process, or even of iterations of an instruction within a single loop. It's a difference of granularity: Multiprocessing runs the spectrum from coarse to fine-grained; parallel processing implies fine-grained.

Most multiprocessing machines can run standard software, but to really take advantage of the power of multiprocessing, you need appropriate systems software. In "A Fearful Symmetry," Craig Keating discusses MPX, a multiprocessor extension to SCO Unix that enables you to get multiprocessor performance from standard Unix applications without modifying them.
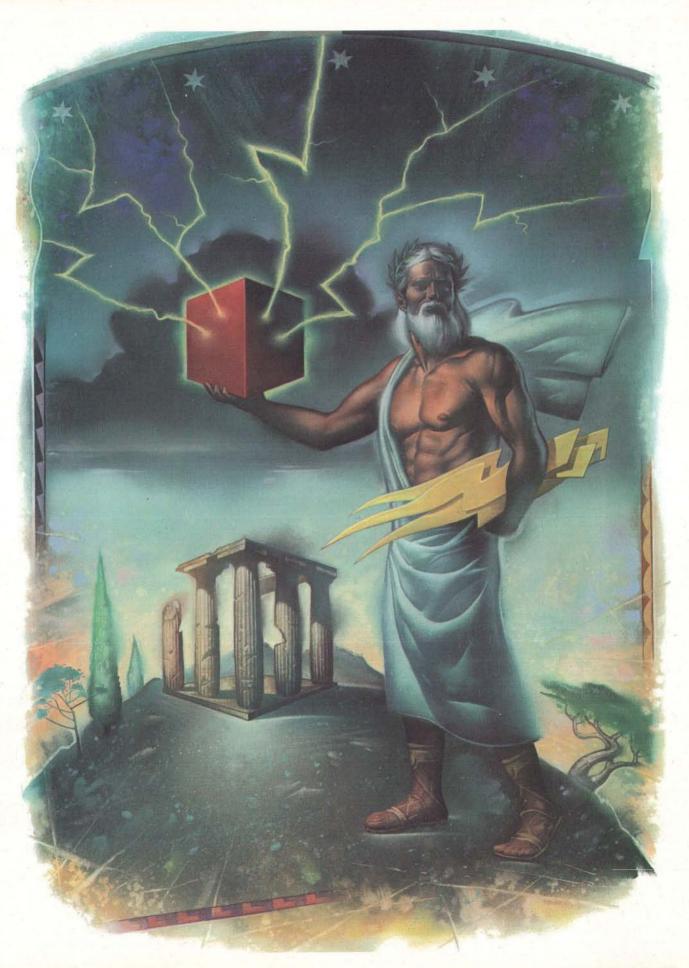
Parallel processing is another technique microcomputers are borrowing from supercomputers, as is specialization. You've heard of dedicated file servers on a network; now there is a dedicated computational server. In "Supercomputers Get Personal," Sam Bogoch, Iain Bason, Jeff Williams, and Mike Russell explore this concept in a specific application: the new ComputeServer from Torque Computers, which uses i860s for parallel processing.

Another aspect of specialization is dedicating certain processors to certain jobs. One of the most CPU-draining jobs is handling I/O. Like supercomputers, in recent years microcomputers have begun to pack this task off to the new intelligent bus masters, leaving the CPU for more "important" work. In "Join the EISA Evolution," Min-Hur Whang and Joe Kua describe the EISA bus master.

Dedicated coprocessors can also be useful. Math coprocessors are essential for heavy-duty number-crunching applications. Graphics coprocessors are becoming the norm rather than the exception. And in "A Calculating RISC," Trevor Marshall tells how RISC coprocessors will soon bring supercomputing power to the desktop.

Perhaps today in the research labs at IBM or Cray or DEC, new hires are looking at rooms full of supercomputers and being told that in 10 years, that power will be available in a 1-foot cube. This time, I'd believe those scientists. But then, this is just hypothetical. Or is it?

*—Jane Morrill Tazelaar
Senior Editor, State of the Art*

Recording the Past...
...Plotting the Future

Our reputation precedes us! From 5 subsidiaries and 35 distributors in more than 40 countries worldwide, thousands of customers purchased more in 1989 than ever before. And they were able to choose new products from an ever-expanding array of plotters, penless plotters, digitizers, recorders and supplies.

The Graphtec reputation is one of building products that work well and last a long time. We earned that reputation the hard way, by delivering over 40 years of the best innovation, support, and after-sales service in the industry.

One of the finest examples of Graphtec innovation is incorporated into our new TM1110/1210 series of thermal plotters. These plotters are unique thanks to the world's first single-substrate, large-format thermal head. The single head ensures a uniform, high-quality printout with superb 406 dpi (16 dots/mm) resolution.

Hard copies can be made on three types of media: regular, thermal transfer and clear film. Also, RS-232C and Centronics interfaces come standard, as do GP-GL and HP-GL™ command sets.

TM1110 A1-size Thermal Plotter

GRAPHTEC

Separated at Birth

*The design of your desktop computer is beginning to show
more than a passing resemblance to that of a supercomputer*

Bob Ryan

I t's funny how things work out. At the same time that Steve Wozniak and Steve Jobs were building Apple I boards in a garage, Cray Research was putting the finishing touches on the Cray 1. At that time, the Apple I and the Cray 1 demarcated the lower and upper bounds of computing power and sophistication. Today, the inheritors of the Apple I tradition are making inroads into areas once considered the sole preserve of supercomputers.

Simply put, supercomputers are the most powerful computers in the world. By using a number of sophisticated design techniques, the fastest supercomputers are able to maintain a sustained throughput of over 1 billion floating-point operations per second (1 GFLOPS) while performing real-world applications. By contrast, the shiny new 25-MHz 486 machine on your desktop can sustain about 1 million floating-point operations per second (1 MFLOPS), or one one-thousandth the throughput of a supercomputer. However, desktop machines are closing the gap, primarily by incorporating some of the same architectural features used in supercomputers.

Making a Computer Super

Computer designers know lots of ways to boost throughput on a computer. The most obvious is to boost the clock speed. For example, the Cray Y-MP has a clock period—the amount of time it takes to execute one cycle—of 6.3 nanoseconds. HNSX Supercomputers' SX-X, due to ship this year, reports a clock period of 2.9 ns. In contrast, a 25-MHz 486 system has a clock period of 40 ns. Obviously, if you could get the 486 to run at supercomputing clock speeds, you would get 10 times its current performance.

Another major advantage supercomputers enjoy over personal computers is pipelining. Successful supercomputer designs ensure that the processor spends most of its time processing. The processor prefetches instructions and data to ensure that it is never idle. Part of the pipelining concept is embodied by intelligent I/O. In a supercomputer, I/O is handled by either a dedicated processor or a separate computer. The supercomputer's arithmetic processors are reserved exclusively for crunching data.

Another important supercomputing concept is matching processor design to the task at hand. Most supercomputing applications involve vector processing, which performs one or more operations on many different pieces of data at the same time. Consequently, supercomputers are optimized to handle vector problems, although they aren't bad at scalar processing, either.

One last feature of supercomputers is *continued*

parallelism—both macro- and microparallelism. (In fact, supercomputing is becoming synonymous with parallel processing.) For example, the SX-X has up to four arithmetic processors sharing a common memory space—an example of macroparallelism. These processors consist of a vector and a RISC scalar unit, with each vector unit consisting of one, two, or four sets of vector pipelines. Each set consists of two add/shift pipelines and two multiply/logical pipelines. Because the pipelines operate simultaneously—an example of microparallelism—each SX-X processor can execute up to 16 FLOPS. That's a theoretical maximum of 5.5 GFLOPS per arithmetic coprocessor. No wonder they call these machines super.

PCs in the Balance
When first introduced, desktop computers were primitive by supercomputing standards. They used slow clocks (the Apple II had a clock period of nearly 900 ns), employed no pipelining or intelligent I/O, enjoyed none of the benefits of specialized coprocessors, and exhibited neither macro- nor microparallelism.

The situation today is far removed from what it was in 1976. The differences between desktop machines and supercomputers are more a matter of degree than a matter of type. The design of most desktop computers exhibits many of the properties that you associate with supercomputers. The increased incorporation of supercomputing features into desktop machines will continue to narrow the gap between desktop and supercomputing performance.

Peddling Faster
Take clock speed. Intel has already begun shipping 33-MHz (30.3-ns clock period) versions of the i486. By the year 2000, Intel expects to be shipping a processor with a 4-ns clock that will still be compatible with the i486. Disregarding all the architectural improvements you can expect in such a chip, the decrease in clock period alone will result in a tenfold increase in speed. Supercomputers can't look forward to a corresponding decrease in clock period. They are already closely bound by the speed of light, which travels about 1 foot in 1 ns.

In the near future, however, it will be advances in architecture, not faster clocks, that bring supercomputing performance to desktop machines.

The Pipes Are Calling
Pipelining—always a feature on supercomputers—is becoming indispensable

M ost PCs exhibit many properties that you associate with supercomputers.

to desktop machines as well. One of the most important features of the i486, for instance, is a 5-stage instruction pipeline that helps keep the processor consistently busy. The new Motorola 68040 goes even further. It maintains separate instruction and data pipelines inside the processor, with separate caches for each.

Although the complex-instruction-set-computer (CISC) processors are catching up, pipelining has had its biggest impact in RISC processors. These usually feature two or more processing units, each with its own pipeline. The units execute in parallel, allowing RISC machines to execute multiple instructions simultaneously. Some RISC processors also feature separate external buses for data and instructions.

Another area where desktop machines are finally emulating supercomputers is in intelligent I/O. Until recent years, most desktop machines used the central processor to move instructions and data between main memory and peripheral devices. Recently, however, sophisticated I/O buses such as the Micro Channel and Expanded Industry Standard Architecture have freed central processors from having to worry about moving bytes around the system. The processor is able to concentrate on processing.

Sharing the Burden
Supercomputers have always featured specialized processors for specialized tasks. Likewise, desktop computers have used specialized coprocessors for years. For example, math coprocessors are available for all popular Intel and Motorola CPUs. The difference is that these processors are now being built onto the microprocessor chip itself. The performance advantages of this integration—as evidenced by the i486 and the 68040—are substantial.

In addition, many companies offer general-purpose and special-purpose coprocessor boards for different desktop machines. These boards perform tasks faster than the host CPU, thus speeding up the system. Some coprocessors offer

better overall performance than the host and simply take over the system. Others, such as array processors, are designed to perform a specific task much faster than the host.

Side by Side
Advances in clock speed, pipelining, intelligent I/O, and coprocessing will all substantially increase the performance of desktop machines in the near future, but parallel processing will surpass all of them in importance. Supercomputers feature both micro- and macroparallelism. Both forms are also showing up on desktop machines.

Microparallelism is the presence of multiple execution pathways within the same processor. For example, the Intel i860 contains three separate pathways—integer, floating-point addition, and floating-point multiplication—that perform math functions. Each pathway has its own pipeline, and they can execute simultaneously. In addition, the floating-point adder and multiplier can be combined to form a vector processor. The i860 represents a trend toward multiple independent execution pathways in a processor, allowing the execution of more than one instruction per clock cycle—a type of architecture called *superscalar*.

As the name implies, macroparallelism does not take place within a single processor, but involves tying together multiple processors of the same type. If the individual processors also feature microparallelism, so much the better. Processors in a macroparallel machine are tied together in two different ways, with a common bus and with point-to-point connections.
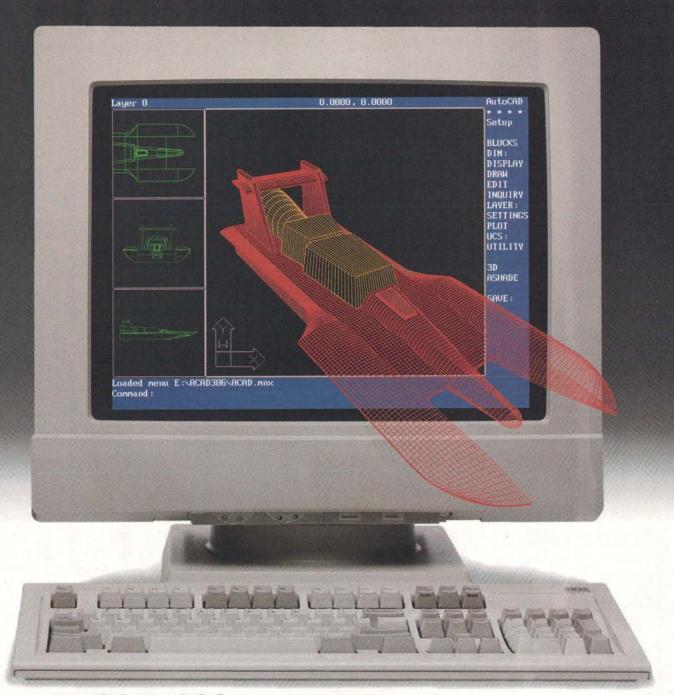
Magic Bus
When connected by a common bus, processors communicate by using shared memory; such machines are thus called shared-memory machines. A good example is the Cray Y-MP, which allows up to eight processors to access the same memory space.

Recently, several companies have produced shared-memory machines based on microprocessors. The Zenith Z-1000 and the Compaq Systempro are the two most notable examples. While shared-memory machines are relatively easy to program, they do have limitations. The biggest is that performance doesn't increase linearly as processors are added.

The more processors there are in a shared-memory system, the more time they spend contending for the bus and keeping data coherent, and the less time

continued

HOW TO COMBINE SPEED, POWER AND GRACE WITH BRUTE FORCE.

AUTOCAD® 386. Zoom faster. Pan faster. Draw faster. AutoCAD 386 combines world-standard CAD performance with full-force 32-bit workstation power—right on your 80386®-based PC.

AutoCAD 386 is built for speed. It loads, redraws, and saves drawings up to 62% faster than before. It accesses up to 16 megabytes of RAM and 4 gigabytes of virtual memory, making more room available for larger AutoLISP® and other applications programs, memory-resident drivers, network interfaces and other utilities—which translates directly into more speed.

Quick! Call Autodesk now to arrange a power-demo at the authorized Autodesk reseller nearest you. And upgrade your present version of AutoCAD for as little as $300.

AUTODESK

2320 Marinship Way, Sausalito, CA 94965
800-445-5415 Ext. 65

Autodesk, the Autodesk logo, AutoCAD and AutoLISP are registered in the U.S. Patent and Trademark Office by Autodesk, Inc. 80386 is a registered trademark of Intel Corporation.

Circle 22 on Reader Service Card

they have to actually process data. As a result, such systems can only scale up so far before the addition of more processors would simply saturate the bus.

Message in a Bottle

The second form of macroparallelism is called message passing. Here, processors don't share a common memory space; rather, they use their own memory and communicate via messages—usually sent over serial links. Arbitration and coherency are not a problem with such systems, but synchronization is. If one processor needs data from another, and the first hasn't finished its computations yet, the second processor must wait until the first processor finishes. The advantage of message-passing systems is that they can theoretically accept an infinite number of additional processors and still obtain linear or near-linear performance increases with each.

Although they are harder to program than shared-memory systems, message-passing systems are showing up both on desktops and in supercomputing centers. Perhaps the most famous example of a message-passing machine is the Connec-

tion Machine from Thinking Machines. This massively parallel machine features over 64,000 identical processors. Recently, Intel announced a message-passing machine—the iSPC/860—that uses from 16 to 128 i860s and can deliver up to 7.6 GFLOPS.

On the desktop, message-passing parallel systems are best represented by transputer-based coprocessing cards. The Inmos transputer is designed specifically as a message-passing processor. Its low cost makes it a natural for desktop systems.

The best thing about parallel-processing technology is that its effects are cumulative. If you increase the clock speed of the individual processors, you increase the speed of the system. Improved pipelining and better use of coprocessors within the individual processors also give a nice boost to system performance, regardless of any improvement to the parallel architecture itself.

Cray Me a River

With most of what makes a supercomputer "super" showing up in desktop machines, you can expect to be able to

perform more and more supercomputing tasks on your desktop. Actually catching up with the Crays of the world is a more difficult proposition. To supercomputer buyers, the cost is secondary to the performance. Thus, supercomputer manufacturers can use the latest technology, no matter what the cost, to achieve ever-increasing performance.

The Cray Y-MP, for instance, uses high-speed, expensive ECL chips for main memory, instead of the MOS chips used in desktop machines. The SX-X uses BiCMOS static RAM. Both use high-speed bipolar elements throughout. Who will be the first to produce a machine based entirely on gallium-arsenide technology? It will be a supercomputer manufacturer; you can count on it.

Although you'll certainly be able to perform many functions once reserved for supercomputers on your desktop, matching the performance of the latest and greatest supercomputer is currently prohibitively expensive. But don't be discouraged; time is on your side. ■

Bob Ryan is a BYTE technical editor. He can be reached on BIX as "b.ryan."

QNX®
The OS for over-achievers™

QNX programmers have a decided advantage.

You see, people who use QNX enjoy the freedom that comes only with a flexible, modular OS. They appreciate the elegance of a **message-passing architecture**. And they marvel at the fact that QNX runs so lean—under 150K—yet out-performs any other PC operating system.

QNX users never worry about whether their applications will make it at runtime, because they know QNX has proven itself again and again in the real world.

It's no wonder that QNX users have achieved so much since the product was first released for the PC in 1982: over 80,000 systems installed in 47 countries world-wide, in all kinds of applications—from making cars to selling books to handling online credit card transactions.

One reviewer dubbed QNX "The multi-everything OS." Now, you might expect

multiuser and multitasking, but realtime? *And* integrated networking? *And* true distributed processing? Best of all, these terms take on a new meaning with QNX.

Multiuser, for instance, means up to 32 terminals per micro. **Multitasking** cashes out as 150 tasks per machine. **Realtime** means not only priority-driven, preemptive task scheduling, but also speed: at 6,896 task switches/sec on a 16MHz 286, QNX is at least a full order of magnitude faster than a typical UNIX system. **Integrated networking** means you won't need yet another layer of software to set up a LAN, and you can use *any mix* of Intel-based micros—from vintage '81 PCs to PS/2s.

Distributed processing with QNX sounds too good to be true. But it is: *Any task can access any resource*—programs, files, devices, even CPUs—without going through the bottleneck of a central file server.

Besides the satisfaction that QNX developers get from using a fast, powerful, and flexible OS, did we mention that they also enjoy *free technical support?*

If you're wondering why you don't already know all about this great OS, you could try asking the over-achievers who are smugly guarding the secret of their success.

Better yet, give us a call. We'll tell you everything you need to know to become an over-achiever yourself.

For more information or a free demo disk, please phone (613) 591-0931.

Circle 223 on Reader Service Card

Spending Your Free Time

A parallel system of supercomputer power is free for the taking in the idle machine cycles of a network's processors

David Gelernter and James Philbin

One way to solve a difficult problem fast is to break the problem into separate pieces and work on all the pieces at the same time. This idea is the whole basis for parallel computing. A parallel program breaks a difficult problem into pieces and then distributes them to a crowd of waiting processors. All the processors work together to find a solution.

Parallelism is a tested and proven way to speed up the process of solving difficult problems. While we used to hear claims that it would be hard to find programs that could be "parallelized," in fact, it has been harder to locate large programs that *cannot* be. Parallel programs are being used in areas such as numerical problem solving, graphics, databases, simulation of large physical and engineering systems, and financial modeling. Today, if you're interested in running programs quickly, you're interested in parallelism, whether you know it or not.

A parallel computer is a single machine that can bring multiple processors to bear on one or more tasks. There are many parallel computers on the market, ranging from workstations with a small handful of processors to huge supercomputers with thousands of processors. Parallel computers are powerful, tend to be relatively simple, and will be a dominating factor in computer architecture in the 1990s.

Parallelism's most important arena, however, may have nothing to do with these machines. With appropriate software and a network, your existing computers can act as a powerful parallel computer.

Hyperactive

To run a program in parallel, you need at least two processors, preferably more, connected together. A parallel computer meets these requirements, but so does a computer network. Whether the processors occupy one box or are strung out across the country, you can take a tough problem, break it into pieces, and hand one piece to each of n processors.

Take a medium-size organization, for example, with a network of 100 workstations or personal computers. How busy are the machines? What portion of processor cycles is actually idle? Assume that during normal work hours, the machines are idle, on the average, about half the time. The real amount of idle time may be a lot higher, however. Even if you run a lot of compute-intensive problems, you also read your mail, edit files, think, eat lunch, and so on. But assuming an average idle time of 50 percent, this organization also owns, on average, a 50-processor (or 50-node) parallel computer, and that's during the workday. Overnight and

continued

on weekends, it might own a 90- to 100-node parallel computer.

How can you use this essentially free parallel computer? When some node on the network has a parallel program to execute, it can offer pieces of the program to other nodes. Some of them will be idle, accept a piece, and begin work. When they are finished, if they are still idle, they will accept another piece and keep going. Some nodes will be too busy to accept a piece. Others will be fairly busy, accept a piece, and progress slowly on it. Through this processor sharing, you can take advantage of quite a bit of the network's unused capacity.

We refer to this kind of network-based parallel computer as a *hypercomputer*. A hypercomputer isn't merely a computer network by another name; it's the parallel computer that *emerges from* the idle portion of the network, like a genie from a bottle. The hypercomputer grows and shrinks depending on the network's unused capacity. It doesn't live in a single box; it takes up temporary residence inside all the network nodes that happen to be available at a given time. You can describe the hypercomputer in statistical

terms: its average power and its average size. Whenever the network expands or machines are upgraded, the hypercomputer's power pool grows.

Processors inside a single-box parallel computer can usually communicate at much greater speeds than networked processors can, but the principle of parallelism is the same in both cases. Because of the difference in communications speed, some programs that run well on parallel computers do not run well on networks. But there appear to be many applications that succeed in both environments.

While hypercomputers can be built on many different network protocols, you need appropriate software to release the genie from the bottle. Hypercomputer Linda, from Scientific Computing Associates, is one example of such software. It is built on top of the Network Computer System/Remote Procedure Call protocol. NCS/RPC, from the Apollo division of Hewlett-Packard, takes care of data conversion between heterogeneous machines. It also takes care of protocol conversion between different network architectures. It runs on LANs and wide-area networks (WANs) and can encompass all

sorts of machine architectures. (Benchmarks of hypercomputer Linda are presented in reference 1.)

Getting Coordinated

Ordinary programming languages aren't geared for writing parallel programs; you need a parallel language or programming environment. A good parallel-programming environment should augment, not obsolete, your favorite languages.

Linda is a coordination language. It supplies the glue that enables you to cement a bunch of independent processes together into a single parallel program. The separate computations can be expressed in any language. A coordination language (e.g., Linda) can be combined with a computation language (e.g., C) to produce a complete parallel-programming language (in this case, C-Linda, also from Scientific Computing Associates). C-Linda systems exist for a variety of shared-memory and distributed-memory parallel systems.

A good parallel-programming environment exists at the language level, not system level. A coordination language

Meet Linda

L inda is an elegant and easy-to-use parallel language. Its power and expressivity result from its use of *tuple space* for interprocess communications and synchronization. Tuple space is a global, associative, object memory that holds *tuples*. A tuple is a sequence of typed fields, where each field is an expression that has a value or a *potential* value. For example,

```
("person", 23, ?x)
```

is a tuple with three fields. Two are values—the string "person" and the integer 23; one is a potential value—the variable x.

Tuples can be inserted into, read or removed from, and evaluated (computed) in tuple space. Four operations accomplish these functions.

```
out(f₁, ..., fₙ)   adds a tuple
in(f₁, ..., fₙ)    removes a tuple
rd(f₁, ..., fₙ)    copies a tuple
eval(f₁, ..., fₙ)  evaluates a tuple
```

The argument to each operation is a tuple, composed of fields f_1 through f_n. Each operation is atomic, so many different programs can safely manipulate tuple space simultaneously.

Adding to Tuple Space
The out operation adds a tuple to tuple space. The value of each tuple field is computed before adding it to tuple space. These are called passive tuples. The eval operation behaves like out, except that field values are determined *after* adding the tuple to tuple space; they are computed concurrently. These are called active tuples. An active tuple turns into a passive tuple when the values of its field expressions have been determined.

Typically, the fields of an eval oper-

ation contain procedure calls. For example, if sqrt(n) is a procedure that yields the square root of *n*, then

```
eval(sqrt(4), sqrt(9), sqrt(16))
```

creates an active tuple with three field expressions and adds it to tuple space. Once in tuple space, the active tuple creates a new thread of control (a process) for each field and computes its value in that thread.

The active tuple above turns into the following passive tuple once all its field values have been determined:

```
(2, 3, 4)
```

Computing a field value in an active tuple may invoke other Linda operations, including eval, to communicate with other processes.

Retrieving from Tuple Space
The in and rd operations retrieve tuples from tuple space. The in operation removes a tuple; rd copies one. Tuples have no names or addresses; rather, in and rd use *associative matching* to locate them. This is similar to the kind of matching that occurs in a relational database.

The arguments to the in and rd operations form a *tuple template*. Templates are formed in the same way that tuples are formed and have the same syntax, but they aren't put into tuple space. Instead, they're used to get tuples out of tuple space. Only passive tuples can match a template; thus, in and rd never obtain active tuples. If no matching tuple is available, in and rd will wait until one shows up.

A field in a tuple can contain a value or a potential value. Potential values have a "?" in front of their field expressions. For example,

```
in("account", id, ?name, ?amount)
```

is an in operation with a tuple template containing four fields, two with values (the string "account" and the variable id) and two with potential values (the variables name and amount). When this template matches a tuple in tuple space, the values of its third and fourth fields will be assigned to the variables name and amount.

Computer Dating with Tuples
If a tuple and a template have the same numbers of fields and their corresponding fields match, then the tuple and the template also match. For a tuple field and a template field to match, the fields must have the same type, and either both fields contain equivalent values or one field contains a value while the other is a potential value.

For example, the tuples

```
("example", i)
( ?s , 3)
(example , ?j)
```

match each other if the variables have the following declarations:

```
int i = 3;
int j;
char example[] = "example";
char s[8];
```

Programs that are designed to use tuple space are portable across widely divergent parallel architectures, including shared-memory, message-passing, and network-connected systems.

FURTHER READING
C-Linda Users Guide and Reference Manual, Scientific Computing Associates, Inc. (246 Church St., Suite 307, New Haven, CT 06510).

supports process creation and interprocess communications. Both services are usually available directly from the operating system in one form or another. Thus, you could bypass coordination languages and rely directly on the operating system for the parallel-programming support that you need. This strategy is simple; unfortunately, however, it is also primitive.

A coordination language is supported

by a compiler and a program-development environment; an operating-system library is not. (The C-Linda compiler, for example, is the most important piece of the C-Linda system. It is implemented as a precompiler generating standard C language code.)

The C-Linda compiler supports clean syntax, compile-time error checking, and high-level run-time debugging and visualization. Most important, it is an

optimizing compiler, so it can support Linda operations efficiently with minimum run-time expense. A communications model like Linda's would be impossible without compile-time optimization. Even a first-rate distributed operating system like Mach is no more a substitute for Linda than it is a substitute for C. Mach and Linda have different purposes. (On the other hand, Mach is an excellent foundation for Linda.)

A good parallel-programming environment must be portable. It must run on any reasonable parallel computer and on any collection of network hosts. Most important, it must be easy to use. Where hypercomputers are concerned, this entails some special requirements.

A Harmonious Match

Parallel programs for hypercomputers face some special conditions. For one, you don't know which processors a given program will run on. For another, you don't know *how many* processors it will run on—a program must be able to take

A tuple space is an object that exists, but you can't say exactly where.

advantage of whatever processors are available at the moment. In addition, network nodes may enter and leave a computation as it progresses.

One approach to handling these conditions is to use the "master/worker" programming paradigm. A master process breaks a program into pieces, or tasks. Each of a collection of identical worker processes grabs a task, performs it, sends the result back to the master process, grabs another, and so on. When no tasks remain, the job is finished.

To build this kind of program with Linda (see the text box "Meet Linda" at left), tasks are dropped into Linda's *tuple space*, one task per tuple. When a worker process needs a task, it simply reaches into tuple space and grabs one. (Tasks can be started in a particular order, if necessary, by using an appropriate data structure. See reference 2.) Tuple space can hold input data or intermediate results as well. When a worker process finishes a task, it drops the result into tuple space, and the master process scoops it out.

This kind of Linda program works the same way regardless of where the worker processes are, or how many there are (as long as there is at least one). Tasks are scheduled dynamically. As a result, the program works the same way with one worker or 100 (100 workers merely run through the pile of tasks faster). Reason-

able load-balancing occurs automatically—one worker process may complete a dozen short tasks while another works on a single long one.

Joining and leaving computations become simple, in principle. A node that was busy but is momentarily free simply grabs a task and begins work on it. A worker process can easily leave between tasks. (Although a worker cannot formally resign from a computation, higher-priority local demands can slow its progress. To cope with this possibility, the master process might reissue tasks under some circumstances.)

Consider the interesting conceptual harmony between hypercomputers and tuple spaces. A tuple space is a kind of shared memory, but unlike other shared memories, the objects in a tuple space have no addresses—they are retrieved by relational database-style matching. So, the physical location and the relative ordering of the tuples are irrelevant.

A tuple space is merely a flock of free-floating tuples; they have no physical addresses to tie them down. A tuple space as a whole is a physical object that exists, but you can't say exactly where.

This same freedom holds for hypercomputers. They exist and can solve problems, but they aren't tied to any particular location, any particular processor. As such, they are in perfect harmony with Linda.

As Invisible as Possible

In a hypercomputer, a node can be anything from a single-processor PC to a shared- or disjoint-memory parallel supercomputer. We'll use the term *node* to refer to any machine connected to the network.

The Linda-based hypercomputer had two principal design goals:

1. It should be invisible to network users who are not hypercomputer users; it shouldn't affect the efficiency or responsiveness of any network node.
2. It should recover as many idle machine cycles on the network computers as possible without compromising the first goal.

The principle of invisibility is very important. You don't want to donate your unused machine cycles to the hypercomputer if your local environment suffers. In practice, you can't achieve total invisibility, but you can come close.

There are three ways in which the hypercomputer can degrade the performance of a network node:

1. It can grab processor cycles that you could have used locally.
2. It can overload the network with hypercomputer traffic.
3. It can force nonhypercomputer programs out of main memory.

If each node can decide when and on what basis it is available, the hypercomputer cannot grab locally needed cycles. A node's local owner decides on an availability policy, and the node announces it to the hypercomputer. Furthermore, hypercomputer processes are scheduled with background priority; therefore, they run only when the node has nothing else to do.

We don't believe that hypercomputer-generated network traffic will be a serious problem. Most networks operate with large reserves of communications capacity. These reserves will be even larger with high-bandwidth fiber-optic networks.

Even in high-traffic networks, communications demands tend to be high only during specific periods of the day (e.g., during network backups). In addition, hypercomputer programs tend to be coarse-grained, to do as much computing as possible between communications operations. Excessive hypercomputer-generated communications affects both the users and the hypercomputer applications adversely, so applications developers will probably avoid it.

Forcing nonhypercomputer programs out of main memory isn't a factor for idle nodes. But for nodes in interactive use, this effect may be noticeable, particularly when you stop interacting with the computer for several minutes and then start again. If you notice this effect and don't like it, you can change the availability policy for your node.

Getting Down to Basics

Configuring the hypercomputer requires listing all member nodes and describing the hardware configuration and availability policy of each. It is becoming increasingly common for LANs and WANs to be interconnected with other networks. Where this is so, nodes must be associated with the hypercomputer they belong to. A node can belong to only one hypercomputer at a time (although it is possible for a node to be part of one hypercomputer during the day and another at night). The hypercomputer configuration file specifies which nodes are potential members of the hypercomputer.

After establishing the configuration, the local-node configuration file must be set up; each potential hypercomputer

node needs its own copy of this file. It has two parts: the local node's hardware description and its availability policy.

The hardware description contains entries such as processor genus and species (e.g., genus Motorola MC68000, species MC68030), number of processors, MIPS (millions of instructions per second) rating for each processor, presence or absence of floating-point accelerators, and (if they are present) their MFLOPS (million floating-point operations per second) rating. The hypercomputer knows this information for most common processors.

The availability policy specifies when the node is available to the hypercomputer. The options include the following:

- time period when the node is available;
- load average above which the node becomes unavailable;
- whether or not the node is available if someone is logged onto it, and if so, after how many seconds of no interaction (this option is mainly for PCs and workstations); and
- maximum number of hypercomputer processes allowed on this node.

If any of the availability criteria is unsatisfied, the node is unavailable.

Daemon Run

The hypercomputer needs a way to distribute the processes among its various processors. The hypercomputer *daemon* (a process that runs continuously in the background) performs this distribution; it has two parts—the scheduler and the allocator. Every node in the hypercomputer runs the daemon. When it starts up, the daemon loads the configuration file to determine which nodes belong to the hypercomputer. If a particular node is a designated scheduler node, the daemon checks with the network location broker. Then the scheduler part of the daemon sits back and waits for nodes to declare themselves available.

The scheduler must be running for the hypercomputer to be available. If the network is up, the hypercomputer can and probably should be available. To ensure this level of reliability, it's a good idea to replicate the scheduler so it can run on more than one node. But it's not necessary; the scheduler can run as a single, unreplicated server. If it is replicated, each scheduler knows where its peer schedulers reside. They communicate among themselves, ensuring that each has an up-to-date copy of the available

processor database. If one of the schedulers fails, the daemon will start a peer to replace it.

The allocator part of the daemon reads the local-node configuration file on start-up, then queries the location broker for the scheduler's address, and registers its node with the scheduler. The registration message includes the node's address and its configuration information. At this point, the scheduler knows that the node is part of the hypercomputer but is not yet available. Next, the allocator determines whether the node is available, using its local availability policy as a guide. If the node is not available, the allocator becomes inactive for a while. It wakes up every so often to keep checking availability.

If the node is available, its allocator sends a message to the scheduler noting current load average and an availability time-out period. The scheduler makes a new entry in its node queue, which is ordered on load average, with least-loaded nodes first. Local allocators send the scheduler refresher messages every so often as notices of their continued availability. If the scheduler doesn't receive such a message before the end of a time-out period, it assumes that the node in question is no longer available. It will assign no further work to that node until it receives a new availability message.

Handing Out the Work

Linda programs create processes on the hypercomputer with the eval operation. When a Linda program wants to create other processes, a request is sent to the local allocator. The message requests one or more processors and may contain processor-specific information such as processor type or minimum MIPS or MFLOPS rating. The message also gives the name of the program (or programs) to run on the allocated processors.

The allocator forwards the message to the scheduler. The scheduler then selects appropriate processors, marks them unavailable, and sends messages to the allocators on each of the selected processors. These messages request that the allocator create a new process and run the specified program.

The local allocator is responsible for creating any hypercomputer processes that are allocated on the local node. When it receives a creation message from the scheduler, the allocator creates a new process, with a background priority, and runs the specified program in it.

A node can run more than one hypercomputer process at the same time. Before sending the "create process" message, the scheduler marks the allocated node unavailable. After creating the new process, the allocator continues to check its node's availability. If it is still available, the allocator informs the scheduler.

Users access the hypercomputer through its shell. This shell includes commands to run Linda programs, compile them for different processor types, show the status of processes and processors, kill a process on any node, kill a program and all its associated processes, and submit programs to be run at a later time. (The hypercomputer software also includes a hypercomputer tuple scope, which is a graphical interactive debugging and analysis utility for use with Linda programs.)

Uncorking the Bottle

Although parallel computing has usually implied multiple processors in one box, it is no longer restricted to this view. The many processors in a network can be harnessed and used for parallel processing as well. You need appropriate software to access this hypercomputer, but other than that, it is essentially free for the taking by using unused, but available, machine cycles from the various processors on the network.

As networks proliferate in the workplace, hypercomputers can use this wasted power. While communications between processors within a single box may be faster, the concept of parallelism works the same way across the network. Hypercomputer Linda is one example of the kind of software that can access this power and let the genie out of its bottle. ∎

ACKNOWLEDGMENT
Some of the early work on Linda systems at Scientific Computing Associates was supported in part by the National Science Foundation Small Business Innovation Research program. The SBIR program is designed to help small firms bring innovative technology to the marketplace.

REFERENCES
1. Philbin, J. "Hypercube Linda Benchmarks." Scientific Computing Associates, Inc. SCA-151, March 1990.
2. Carriero, N., and David Gelernter. "How to Write Parallel Programs: A Guide to the Perplexed." *ACM Computer Surveys*, September 1989.

David Gelernter is an associate professor of computer science at Yale University (New Haven, CT). James Philbin is a research scientist and product manager for Scientific Computing Associates, Inc. (New Haven, CT).

A Fearful Symmetry

*Watch out, uniprocessor systems—here's systems software that
can harness the power of multiple processors*

Craig Keating

Multiprocessing
is hot. With
companies like
Compaq and
Zenith Data Systems now
selling systems that harness
the power of multiple micro-
processors, personal com-
puters can compete in areas
formerly reserved for much
larger systems. To comple-
ment their multiprocessing
hardware, some of these com-
panies are turning to MPX
from The Santa Cruz Opera-
tion (SCO) as one means of
providing multiprocessing
software.

MPX—a multiprocessor
extension to SCO Unix—is
based on smp (for symmetric
multiprocessor) from Corol-
lary. It is an off-the-shelf soft-
ware package that runs on
several multiprocessing com-
puters, including the Corol-
lary 386/smp and 486/smp,
the Compaq Systempro, the Zenith Z-
1000, the Mitac Series 500, and the
Apricot MC486. MPX is designed to let
you run software developed for SCO
Unix on a multiprocessing system with-
out modification. It also lets you take ad-
vantage of standard peripheral devices
by running unmodified device drivers. A
look at it points out the problems inherent
in designing multiprocessing systems.

MPX Roots
Corollary has been developing multipro-
cessor software and hardware for over
four years. Its primary goal is to provide
core multiprocessor systems that take ad-
vantage of off-the-shelf third-party soft-
ware and peripherals (see the text box
"The smp Hardware" on page 222). For
this reason, the company discarded soft-
ware solutions that involved message-

passing and proprietary oper-
ating systems. Instead, smp
consists of numerous en-
hancements to an already
popular operating system—
SCO Unix System V/386.
SCO Unix was used as the
base for smp because of its
dominant market position, its
large installed base, and the
many software applications
and device drivers available
for it.

To maintain compatibility
with SCO Unix applications,
Corollary restricted all multi-
processing modifications for
smp to the kernel. The smp
kernel consists of one text seg-
ment and one data segment,
just like the uniprocessor ker-
nel. All CPUs in the system
run a single copy of the kernel
text and share the same kernel
data. The smp kernel runs
standard off-the-shelf Unix
applications.

In addition, the smp kernel supports
AT-bus peripheral devices with off-the-
shelf Unix device drivers. This software
and hardware compatibility doesn't re-
quire any source code changes or recom-
pilation. The fact that an smp system is
running with more than one CPU is
transparent to the user in every way ex-
cept performance.

The smp Hardware

On the hardware side, Corollary provides a core computer system consisting of CPUs, caches, memory boards, and buses. This technology, termed 386/smp and 486/smp, lets you add more than one CPU to a PC architecture. In fact, the first CPU in the system is completely PC compatible; it has access to global shared memory and an AT bus. Since the AT bus is not suitable for multiple CPUs, the smp architecture includes a second bus called the C-bus (cache bus). The AT bus is used for booting the system and for third-party peripherals, while the C-bus provides increased memory bandwidth, cache coherency, and fair arbitration.

To satisfy the memory bandwidth that multiple CPUs require, each CPU comes with a local 256K-byte direct-mapped, write-back cache that transfers written cache data to the common shared memory only when needed. The write-back cache reduces bus and memory contention by approximately a factor of five when compared to a comparable write-through cache. Minimizing bus traffic lets you connect up to 10 CPUs to the shared memory while keeping the bus utilization under 50 percent.

The 486/smp is a global shared memory architecture (see figure A). Each CPU may or may not have I/O interfaces associated with it. Figure A shows two CPUs with serial I/O, one with a SCSI port, and one with an AT bus. This not only prevents I/O bottlenecks by distributing the I/O among different CPUs in the system, but also minimizes the number of slots needed to create a powerful system. Performance-critical I/O (i.e., disk and serial I/O) has been
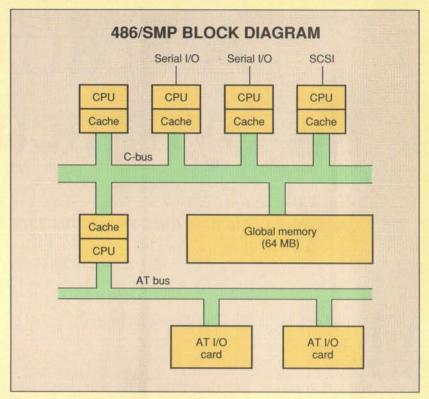
Figure A: *The 486/smp block diagram shows the dual-bus nature of the smp hardware. The default CPU provides access to standard peripherals via its connection to the AT bus. The high-speed C-bus connects the processors to each other and to the shared memory. Optionally, the processors can also provide direct I/O to the C-bus.*

moved to the additional CPUs. This utilizes the significant bus-speed advantage of the C-bus while freeing up the scarce bandwidth of the AT bus. The SCSI disk adapter, in particular, was designed with scatter-gathering DMA logic and specialized to take full advantage of the C-bus.

The smp hardware is used in the Zenith Data Systems Z-1000 and Mitac Series 500, as well as in Corollary's 386/smp and 486/smp.

Shadowed Structures

Because existing Unix utilities and device drivers access routines and data structures in the kernel and rely on their results or contents, certain parts of it could not be modified for smp. To maintain full compatibility with all Unix applications and device drivers, the file system layout, system call interface, object and executable file format, kernel data structures, and kernel routine parameters, actions, and return values remain unchanged. However, the kernel must maintain more information than the standard Unix data structures contain to allow multiple CPUs to run tasks simultaneously. This additional information is kept in new data structures established in parallel to the standard ones; this maintains both applications software and driver compatibility.

You install the smp kernel by installing new object files into the Unix Link Kit and rebuilding the kernel. Booting the smp kernel is the same as booting a standard Unix kernel on a 386 or 486 machine: The BIOS ROM loads the kernel into memory and passes control to it. Once booted, the smp kernel determines the multiprocessor-platform type (e.g., 486/smp, Systempro, or Z-1000) and the number of CPUs in the system.

The smp kernel was designed to meet a number of important goals. It must provide both source code and binary compatibility with existing Unix applications and device drivers. It must also require minimal changes to the uniprocessor Unix kernel, take an incremental approach to multiprocessing modifications, provide hardware portability, and result in increased performance under intensive work loads as the number of CPUs increases from one to 16.

Kernel modifications for smp were made strictly to support multiprocessor technology. For the most part, these

changes involved making areas in the kernel reentrant, not rewriting large portions of it. The minimal changes involved in producing reentrant code reduce the risk of incompatibility. In addition, minimizing kernel changes simplifies the task of tracking future Unix releases.

Multiprocessor technology greatly increases the number of users the system can handle. An smp system can start out supporting a few users and incrementally grow to support over 200. This additional capability forced some kernel profiling and resulted in optimizations that provide increased performance without compromising compatibility.

The Quest for Compatibility

Compatibility is achieved by restricting modifications made to the kernel and by establishing certain guidelines for those modifications. In keeping with the philosophy that multiprocessing should be transparent to the user, you don't need to modify application programs to take advantage of multiple CPUs. Thus, the smp kernel can run all the standard Xenix/Unix applications available from third-party software vendors.

Maintaining driver compatibility is a more difficult task. To ensure it, kernel header files, routine parameters, and return values are left unaltered. This allows the smp kernel to run device drivers from third-party hardware vendors unmodified. In fact, the device-driver source code doesn't even require recompilation.

The Incremental Approach

In its earliest development stages, the smp kernel was not very symmetric. Only the CPU that booted the system (the "default" CPU) could execute system code; the other CPUs executed user code only. During the design process, some added testing using commercial applications determined the frequency with which each system call was used. Then, on an incremental basis, the system calls that composed 80 percent of the total call frequency were modified so that any CPU could execute them. However, the default CPU must still execute the remaining system calls. The system uses a table—established at link time—to determine (at run time) whether a given system call can execute on any CPU in the system or only on the default CPU.

At first, the smp kernel also restricted device drivers to the default CPU, so you could link any device driver to the kernel without modifying or even recompiling it. However, now you can modify performance-critical device drivers so that any CPU in the system can execute them. (Unmodified drivers will still run only on the default CPU.) The system maintains a table that shows which device drivers have been modified. The smp kernel can thus determine at run time—just before calling a device driver—whether any CPU or only the default CPU can execute its code.

The incremental approach used to modify system code and device drivers points out the advantages of having a default CPU. This CPU provides the hardware and software environment that standard drivers expect; thus, all existing or future device-driver code can execute unmodified. Letting you take advantage of the latest innovations in driver software as it becomes more advanced is an important feature of smp. The default CPU also minimizes engineering risk by letting Corollary and third-party device makers test each modified driver or newly distributed system call.

Scheduling User Processes

Once you boot the smp kernel, the system can execute any user program (i.e., process) on any of the CPUs in the system. The scheduler code chooses the process to be executed next by selecting the one with the highest priority. The scheduler is multithreaded, so all CPUs can execute it simultaneously. Each CPU is thus self-scheduling: Each is responsible for selecting which process it will execute. The alternative would be to have a single controlling CPU using a complicated load-balancing algorithm to migrate processes to the individual processors.

Like its uniprocessor cousin, the smp kernel has a single run queue that all the CPUs in the system share. The run queue is a linked list of all processes that are ready for immediate execution. When a CPU is available, it removes the highest-priority process from the run queue and executes it. When the process is no longer executable (e.g., when it is waiting for I/O), the CPU returns it to the run queue and chooses another process.

The scheduler in the smp kernel uses the standard Unix priority algorithm to select a process. It has only been modified to accommodate the default CPU. When a CPU other than the default encounters unmodified device-driver code or a nonmultithreaded system call, it returns the process to the run queue after marking the process as one that can run only on the default CPU. The scheduler will let only the default CPU select a process that is marked. As soon as the default CPU executes the marked section of code, the process is once again marked as executable by any CPU.

The default CPU chooses marked processes before any others, regardless of the priorities of the unmarked processes. If it finds no marked processes, it will execute the highest-priority unmarked process. The default CPU must search the run queue only once to determine its next process, regardless of whether it chooses a marked or an unmarked process.

The smp system provides the scheduler with additional multiprocessor system-control features through a pseudo-driver. You invoke these additional control features using calls to the driver. Unlike using additional system calls, using a driver doesn't compromise compatibility. The pseudo-driver supports the following additional features:

- enable/disable CPUs from executing processes;
- lock a process to be executed by a specific CPU; and
- force a CPU to execute only processes that are locked on it.

These features were crucial initially for smp kernel development and debugging. Because of this, the system still furnishes both the interface to these driver calls and an application program that uses them. In general, however, the scheduler performs markedly better without using these features.

Getting in Sync

In a multiple CPU environment, processors execute common kernel code and operate on common data structures. The mechanism that guarantees mutual exclusion to code and data is called *synchronization*.

Many sections of code—called critical sections—cannot tolerate more than one CPU executing them at the same time. There are a number of ways to provide mutual exclusion using software-locking mechanisms, but they are costly and impose a lot of extra overhead. The smp kernel requires synchronization support in hardware.
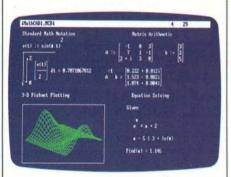
Specifically, the smp kernel requires that the hardware support "test and set." This is typically supported with the xchg instruction in the 386 and i486 processors. These chips have a lock signal that the hardware can use to support indivisible bus cycles. The xchg instruction will assert LOCK for the duration of the instruction; therefore, no other CPU can access memory between the read and the write exchange.

continued

After centuries of practice, mankind perfects engineering calculations: MathCAD.

Announcing MathCAD 2.5: The Dawn of a New Age.

What the historians will call it, only time will tell.

Perhaps the Century of Speed, or the Era of Ease. But whatever the name, this is the age of MathCAD 2.5, the only math package that looks and works the way you think.

MathCAD 2.5 includes 3-D plotting, HPGL sketch import, and PostScript output.

MathCAD is far and away the best-selling math package in the world. Because it lets you perform engineering and scientific calculations in a way that's faster, more natural and less error-prone than the way you're doing them now—whether you're using a scratchpad, calculator, spreadsheet or program that you wrote yourself.

And now there's MathCAD 2.5, a dramatically improved version that includes three-dimensional plotting, enhanced numerical analysis, PostScript® printer support, and HPGL file import from popular CAD programs like AutoCAD.®

And like before, MathCAD's live document interface™ lets you enter equations anywhere on the screen, add text to support your work, and graph the results. Then print your analysis in pre-sentation-quality documents.

It has over 120 commonly used functions built right in, for handling equations and formulas, as well as exponentials, differentials, cubic splines, FFTs and matrices.

No matter what kind of math you do, MathCAD 2.5 has a solution for you. In fact, it's used by over 90,000 engineers and scientists, including electrical, industrial, and mechanical engineers, physicists, biologists and economists.

But don't take our word for it; just ask the experts. PC Magazine recently described MathCAD as "everything you have ever dreamed of in a mathematical toolbox."

And for Macintosh® users, we present MathCAD, rewritten to take full advantage of the Macintosh interface. Entering operators and Greek letters into equations is pure simplicity!
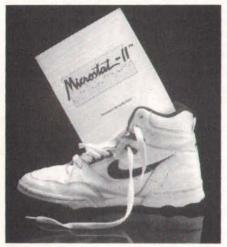
Look for MathCAD 2.5 at your local software dealer, or give us a call. For more information, a free demo disk, or upgrade information, dial 1-800-MATHCAD (in MA, 617-577-1017).

Available for IBM® compatibles and Macintosh computers.
Call for UNIX platform availability.

TM and ® signify manufacturer's trademark or manufacturer's registered trademark respectively.

UNIX version now available.

PC MAGAZINE EDITOR'S CHOICE
March 14, 1989 issue.
Best of '88
Best of '87

MathCAD®

MathSoft, Inc. 201 Broadway, Cambridge, MA 02139

A unique lock protects each critical section in the smp kernel. Each lock is a 2-byte structure denoting busy and own-ership conditions. The busy byte con-tains either a 0 or a 1 (a 1 means that a CPU currently owns the lock). The own-ership byte holds the identification of the owning CPU. Primitive locking and un-locking routines in assembly language use the xchg instruction to acquire locks. A CPU acquires a lock prior to the criti-cal section and releases it after that sec-tion has been executed.

Put a Lock on It

If a lock is busy when a CPU wishes to acquire it, the CPU waits in a spin loop (see figure 1). Corollary performed ex-tensive testing to both discover which were the critical sections and guarantee that they were of short duration. Short critical sections allow the smp kernel to wait for a lock rather than having to per-form an overhead-intensive context switch to another process, thereby com-promising performance. (For more in-formation, refer to *The Design of the Unix Operating System* by Maurice J. Bach [Prentice-Hall, 1986].) Each lock in the smp kernel has been evaluated for its frequency of lock requests and the number of spin loops executed while waiting for it. This testing verifies that spin loops due to locked critical sections occur infrequently.

It is important to note that a CPU ac-quires a lock on behalf of a particular process. Each CPU maintains its own
continued
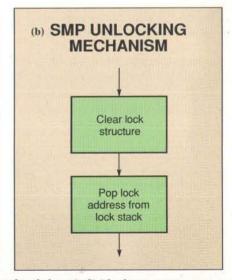
Figure 1: *Locks on critical sections of kernel code keep individual processors from corrupting the common shared kernel data.* **(a)** *Acquiring a lock can involve a spin loop when another processor already owns the lock.* **(b)** *Freeing a lock is straightforward.*

Securing your data is as fast as one, two!

When you use DPT's

Smart Cache™

the caching controller with disk-mirroring.

Speed and security in one package.

SmartCache with disk-mirroring ensures uninterrupted operation in the event of disk failure by simultaneously writing data to your PC's primary drive and a secondary, "mirror" drive. SmartCache increases system speed and eliminates down time and data loss due to drive failure, thus providing true disk fault tolerance for all PCs.

When we say *fast*...

Unlike software disk-mirroring that actually slows your PC down, SmartCache dramatically boosts system speed, thanks to its state-of-the-art disk-caching technology. The on-board processor and expandable cache RAM allow the controller to simultaneously manage the primary drive and the mirror drive and still process data at 0.5ms for unprecedented speed and security.

Picture-perfect compatibility.

SmartCache needs no special ROM BIOS or software drivers and is fully compatible with *all* 286/386 PCs and *all* operating systems. Because SmartCache looks exactly like a standard disk controller to your computer, installation is quick and easy.

So smart, it can repair itself.

SmartCache actually repairs all disk errors *automatically*. Even catastrophic drive failures will not cause the system to crash.

Backed by the best: DPT.

Distributed Processing Technology was the first to develop caching disk controllers for microcomputers and is the recognized leader in the industry. Our products have been at work for over a decade, speeding up minis and mainframes. We offer a 1-year warranty, clear documentation, and outstanding technical support.

Look into SmartCache!

You'll like what you see. Call today and find out how you can add speed and security to your PC system—with SmartCache, from DPT.

SMP PERFORMANCE BENCHMARKS
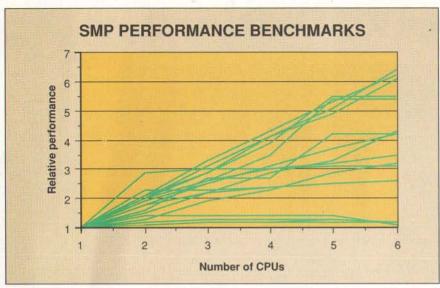
Figure 2: *Graphing the results of the Neal Nelson and Associates Business Benchmark reveals how performance under the smp kernel increases as more CPUs are added to the system. Each line represents a different type of task. Some of the tasks show a near-linear increase. The flat lines represent disk-access time tests; disk access doesn't improve as you add CPUs.*

lock stack so that it knows about all the locks acquired for a particular process. Just as important, being a stack, it also maintains the exact order in which the processor acquired the locks. To avoid system deadlock, some critical sections have an imposed lock hierarchy, and their locks must be acquired in an exact sequence.

When a process in the smp kernel suspends itself by calling sleep(), all the locks it holds at that point must be released or the system could deadlock. Because the lock stack maintains acquired locks, releasing them is a simple procedure: The system copies the current state of the lock stack to the executing process's kernel stack and then releases the locks. This gives other CPUs access to the locks. When the process resumes, the system uses the saved lock stack to reacquire the same locks. The process then continues executing from the critical section where it suspended itself.

The smp kernel doesn't add any new situations that cause a process to suspend itself that aren't in standard Unix. This is another factor that lets the smp kernel work with multiple CPUs without extensive modifications to standard Unix. It also makes tracking future releases of Unix as easy as possible.

You can think of some of the kernel's critical sections as a group of critical sections with small pieces of reentrant code in between. Rather than acquiring and freeing locks on many small sections,

the process acquires a single lock for the group. This allows the smp kernel to minimize the locking overhead in the system. The lock is acquired only once as long as the reentrant portion of the critical section is short.

The lock and unlock routines are the only parts of the kernel that are responsible for acquiring and releasing locks, and for maintaining the lock structure and lock stack. Both routines are coded in assembly language and streamlined for the typical case: lock availability. Timings with a 25-MHz 386 CPU determined that a lock takes 16.6 microseconds and an unlock takes 12.1 μs. On a 25-MHz i486 CPU, these figures drop to 6.6 μs for a lock and 4.8 μs for an unlock.

The Other Half

The smp kernel was originally designed for Corollary's 386/smp and 486/smp products. Currently, however, it also runs on other systems. Ports to other hardware systems took place as manufacturers began searching for multiprocessor systems software. Because the hardware-dependent portions of the smp kernel were segregated into a couple of kernel source files, initial ports typically took weeks.

Since then, the smp kernel has been further modularized to support new multiprocessor platforms without smp kernel source code. A single file of 10 to 20 short hardware-specific routines and an entry in a hardware-platform multi-

processor switch table are all you need. Because of this modularization, ports can now be done in a matter of days.

The smp kernel doesn't exist in a vacuum. No matter who the manufacturer is, the kernel requires that the underlying hardware architecture include four vital features: global shared memory, support for test and set, individual CPU-to-CPU interrupts, and cache transparency.

You need global shared memory because the smp kernel requires that all CPUs have access to all memory. In addition, all CPUs must view memory identically: They must all see the same physical byte of memory at the same location. Also, the smp kernel assumes that there is no "distant memory" for any of the CPUs.

The smp kernel uses test and set as the mechanism that ensures that kernel data structures and critical sections are protected. This capability must be present in the hardware. The smp kernel further requires that each CPU be able to individually interrupt all other CPUs and that this interrupt be maskable to ensure the protection of critical sections of software.

Most multiprocessor architectures feature CPU caching designs to maintain high system performance. The smp requires that the hardware provide for cache coherency and that, other than initial cache flushing and enabling, the coherency scheme be transparent to the kernel.

Compatibility Plus

While the smp kernel was designed with compatibility as the highest concern, performance was not overlooked. As figure 2 demonstrates, performance under a 60-user work load increases significantly as you add CPUs to the system.

The popularity of SCO Xenix and Unix means that a large proportion of Unix implementations use Intel-based platforms. Adding multiprocessor technology provides these machines with a growth path. It gives you the ability to increase performance as your system needs change.

The Corollary smp kernel, as incorporated into SCO MPX, provides a growth path that is compatible with PC platforms and SCO Unix. In addition, the kernel offers a binary-compatible operating-system platform that you can run on powerful multiprocessing platforms. ∎

Craig Keating is manager of multiprocessor software for Corollary, Inc. (Irvine, CA). He can be reached on BIX c/o "editors."

BENDING OVER BACKWARDS WON'T GET YOUR APPLICATIONS UP AND RUNNING

When your front-end tools fall short, don't get bent out of shape. Get JAM, the user interface management system that goes far beyond the front end tools provided by other vendors. JAM has everything you need to build **powerful,** completely **portable** applications **quickly** and **easily.** ▶ With JAM, you can create and link together the screens, windows, fields and menus that make up an application. The result is a working **prototype** that incorporates special features like extensive field edits and validations, context-sensitive help, pop-up windows and pull-down menus—all **without coding.** ▶ And if you're working with databases, JAM's database interface, JAM/DB*i*, links your applications to many popular databases. Access is implemented directly in JYACC's procedural language using SQL. ▶ JAM applications are **fast** because processing routines are written in standard third generation languages. And, you'll never have to worry about converting your applications to a new system, because JAM is hardware, operating system and database **independent.** ▶ So, if your development tools have you turned upside down, it's time you turned to the experts. JAM, by JYACC. Specialists in application development. ▶ To receive a *FREE* demo diskette and find out more about the JAM family of software products, call our toll-free hotline today!

800-458-3313
(In NY, call 212/267-7722 or FAX 212/608-6753)

☑ Supports MS/DOS, OS/2, XENIX, UNIX, VMS, Primos, AOS/VS, VOS, MPE/XL, iRMX

☑ Links to Informix, INGRES, Oracle, Rdb, Share Base, SQLBase, SQL Server, Sybase, xdb

JAM

JYACC

Supercomputers Get Personal

Torque's new i860-based ComputeServer is fresh and fast—
an attractive way to serve up computational power

Sam Bogoch, Iain Bason, Jeff Williams, and Mike Russell

Bringing supercomputer performance to your desktop is not so much a question of technology as it is one of economics. Many companies build workstations and coprocessors that, at least in some areas, deliver supercomputing performance from desktop machines. The problem is that you can spend a fortune trying to outfit everyone in your organization with high-priced hardware. The ComputeServer from Torque Computer offers an affordable alternative.

The ComputeServer (see the photo on page 233) is a parallel-processing computing engine that delivers supercomputing performance by putting a "desktop supercomputer" on a LAN. As a peripheral resource for personal computers and workstations, it is analogous to already-popular print and file servers found on most LANs. In fact, it is the large installed base of standard networks (the ComputeServer uses Ethernet) that makes the computational server concept viable. The server approach is also consistent with the needs of most desktop "power" users who wouldn't use a supercomputer full-time but would require its power in bursts.

Client-Server Computing
The ComputeServer uses a client-server architecture—much like the database servers that are becoming so popular on LANs. Applications are split into client and server portions. The client portion runs on your desktop machine and provides the user interface for the application. The server portion consists of computationally intensive code residing on

the ComputeServer. The client calls this code as needed, and the server returns the results (see figure 1).

While the ComputeServer is transparent to users, it is not transparent to programmers. Applications must be modified to take advantage of its power. Fortunately, many programmers are already using structured-programming techniques to divide user-interface modules from computationally intensive ones, so converting code to support a true client-server approach is not difficult.

The system software takes advantage of modular programs by using the Linda memory model originated by David Gelernter at Yale (see "Getting the Job Done," November 1988 BYTE). Linda adds just six statements to a conventional language such as C, yet it easily splits client and server functions and handles all the communications between the two. Just as important, Torque's Linda implementation allows multiple server subtasks to execute simultaneously on the ComputeServer's multiple processors.

Linda automatically handles the two main problems in getting one or more

COMPUTESERVER SOFTWARE

Desktop system

ComputeServer

Tuple space

Regions to process (based on number of processes), parameters (set, pixel granularity, etc.), control tuples (start, stop, quit)

Completed scanlines (arrays), progress information (errors, out of memory, etc.), ready signals from processes (optional)

Boundary conditions

User interface
eval
out
in

Process
in
out

Process
in
out

Process
in
out

Process
in
out

Figure 1: *A ComputeServer application exists on both the client and server machines. Communication between the two parts occurs in tuple space.*

applications to run on multiple processors: process creation and data consistency. It spins off tasks to run on the different processors and makes certain the tasks send and receive the right data at the right times.

Torque's implementation of the Linda model asynchronously spins off C functions, as well as linked code in FORTRAN 77 or other languages, onto one or more remote processors. Linda treats the multiple CPUs as a processing pool. It can allocate any processor to one or more subtasks that the application programs create.

Linda also creates software-based global memory for passing data between the client machine and the ComputeServer's multiple processors. The system subdivides the available processor pool on the fly among multiple clients and keeps each client's memory space separate during execution.

End of the Rainbow
The processors at the heart of the ComputeServer supply its supercomputing power. The system contains from one to 16 Intel i860s, which offer marked performance advantages over other current RISC and complex-instruction-set-computer processors (see "Intel's Cray-on-a-Chip," May 1989 BYTE).

The 1-million-transistor i860 includes a RISC core, a vector-capable 64-bit FPU, 4K-byte and 8K-byte data caches, and hardware support for three-dimensional graphics primitives. The tight coupling of these functions within a single chip via fast, wide (up to 128 bits) on-chip buses provides much faster floating-point performance than the multichip sets needed to implement other RISC architectures.

The i860 can reach 66 million floating-point operations per second at peak performance, although only a few applications, such as neural-network simulations, can effectively harness all the potential of the microprocessor's multiply-accumulate pipeline. A more realistic performance figure is the 17 MFLOPS quoted by Intel for running an optimized LINPACK suite. The processor should be able to sustain 10 MFLOPS in most nonoptimized applications.

So, a ComputeServer with 16 processing units is capable of more than 1000 MFLOPS at peak performance, and 160 MFLOPS sustained. By comparison, a Mac IIx can sustain 0.4 MFLOPS; a SPARCStation 1, 1.2 MFLOPS; and a Cray X-MP with four processors, 200 MFLOPS. A uniprocessor ComputeServer for $20,000 runs nearly as fast as the original Cray 1 on many applications.

The i860 also offers an upgrade path to superscalar technology, also known as very long instruction word technology (see "VLIW: Heir to RISC?" August

1989 BYTE). A superscalar architecture features multiple integer and floating-point arithmetic units within a single processor. Advanced compilers can simultaneously assign different jobs to these processors. Superscalar is a form of parallel processing called *microparallelism*, which, unlike the *macroparallelism* of the ComputeServer's multiple processors, is transparent to the programmer. As superscalar technology matures, future Torque machines will continue to incorporate both microparallelism and macroparallelism.

Ties That Bind
One of the central concerns in parallel processing is how to most effectively tie multiple processors together. In distributed-memory architectures, each processor has private RAM and communicates with the other processors via messaging links. In shared-memory architectures, processors share the same bus and main memory.

Often, software development issues obscure the debate about the underlying pros and cons of each of these hardware architectures. The key issue becomes which architecture is easier to program for a given job, rather than which one is inherently better. Because programming a shared-memory machine is often easier than programming a message-passing one, most people are willing to live with the higher hardware costs and lack of scalability associated with the former.

However, C-Linda and other new development tools have made these architectures interchangeable from a programmer's point of view. Either architecture, or a hybrid of the two, can perform most parallel-programming tasks handily. Linda has shown that the real key to successful parallel processing is a machine's ability to provide high bandwidth and low latency for interprocessor communications.

The Virtual Tree architecture used in the ComputeServer combines many of the advantages of shared-memory and message-passing architectures. Developed by the ComputeServer design team in 1987 for the multi-8086 Parallon processor, the VT uses a layered hierarchy of messaging buses optimized for burst-mode transmission. The tree is called "virtual" because each branching layer is implemented as a fast-messaging bus rather than as many slower point-to-point links.

Unlike the buses on shared-memory systems, the fast-messaging buses pass messages rather than access and arbitrate shared-memory locations. Thus, you

don't need the complex bus-snooping hardware of shared-memory systems, and processors on one bus do not need knowledge of transactions occurring on other buses.

The kicker is that, because the system passes messages by way of these very fast shared channels, it can adapt to changing communications loads without employing any complex routing overhead. The VT also provides a hardware-broadcast facility that dramatically improves bus utilization when you have to share data among many processors.

The single-layer VT bus used in the ComputeServer is 64 bits wide, features burst mode, and is capable of 66-megabyte-per-second point-to-point transfers and 1-gigabyte-per-second effective broadcast transfers when all 16 processors are "listening." The bus supports up to 32 devices within a given layer. The system can employ addresses that the processors don't use to accommodate parallel I/O devices such as multiple disk drives, frame buffers, and real-time data acquisition systems.

Each processor board contains one or two complete i860 processing units, each with up to 16 MB of local static-column DRAM, and interfaces to both the VT and the system I/O bus (see figure 2). The ComputeServer backplane can hold up to eight processor boards (16 CPUs). This modular architecture allows board-level upgrades to increase the number of processors, or board swaps to new CPUs. For instance, if you have an earlier, transputer-based ComputeServer, you can upgrade to an i860 system for the cost differential between the new system and your current one.

ComputeServer I/O

Central to the computational-server concept is the ComputeServer's ability to connect to standard networks. Although the I/O requirements for a computational server are less complex than those for file and database servers, rapid response to input tasks and sustained network performance are still critical. The system employs a dedicated 386-based processor to handle I/O functions, and it features an Extended Industry Standard Architecture backplane to ensure compatibility with both current and future high-speed networks.

The I/O processor runs Unix System V/386 and manages the ComputeServer's built-in hard disk, which stores both executable code and temporary data files. The widespread availability of networking hardware and drivers for this I/O processor and bus combination
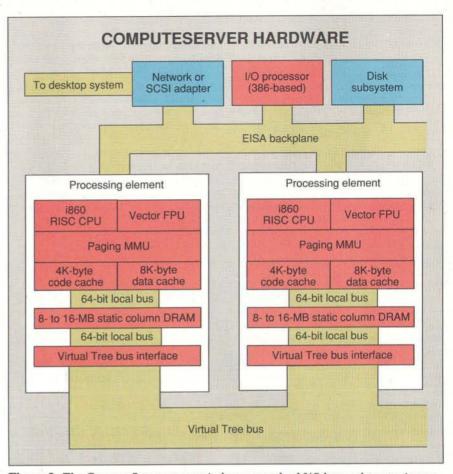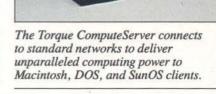
Figure 2: *The ComputeServer uses an industry-standard I/O bus and a proprietary fast-messaging bus. Each processor has its own local memory.*

means that the ComputeServer can support new networks without redesign. The high-level system software must simply be ported to the new I/O boards.

Software Architectures

For all its high-tech hardware, the ComputeServer would be useless without software capable of harnessing its power. One option considered was to implement a system using the X Window System's (referred to as X Window hereafter) smart-terminal model of computation.

The X Window model has a limited set of primitives for the desktop machine (the "window server" in X Window terminology) to execute, and it effectively runs the entire application on the server (the "window client"). These standard primitives express any communication between client and server. This places the desktop system in the role of intelligent terminal; it translates your input into primitives that the server uses, and primitives from the server into on-screen graphical elements.

The Torque ComputeServer connects to standard networks to deliver unparalleled computing power to Macintosh, DOS, and SunOS clients.

One problem with this model is that response time and interactivity suffer because of the communications load on the intervening network. Another is that the user interface is defined by the computational-server "client," not the desktop "server," sometimes forcing you to straddle different styles of computing when you switch from local applications to remote ones.

While the ComputeServer can support X Window applications, it doesn't use the X Window model as its primary mode of communications. A true computational server should fit transparently into desktop environments, rather than the other way around. You should not have to log onto a remote system; your applications should do that for you.

The desktop machine should run the bulk of a server-capable application; this ensures the highest system responsiveness. Only those routines that require sustained computation should be farmed out, on the fly, to the computational server. One of the biggest challenges, therefore, is to provide the tools needed to build this new class of applications.

Implementing Linda

The ComputeServer layers its Linda environment above industry-standard software elements. The I/O processor's use of System V/386 means that you can log onto the system under Unix. In this capacity, the system can be used as an X Window client.

More important, the ComputeServer uses TCP/IP as the lowest internal layer of Linda communications between desktop machines and the ComputeServer. When compiling a client/server application, the system links the appropriate communications stubs into the run-time libraries at each end. The exception to this rule is the SCSI direct-connect option, which lets you connect directly to the ComputeServer, monopolizing all the processing power in the system. This connection uses a more direct transport mechanism.

Client/server applications handle all connectivity between the ComputeServer and desktop machines. You don't have to modify or bypass your desktop computer's system software. You also don't have to install dedicated drivers or update ComputeServer applications when you upgrade to a new version of your desktop operating system.

Sharing, Not Shared, Memory

Linda uses *tuple space* to maintain data consistency among far-flung processes. In effect, tuple space is a database that is shared by processes within a Linda program. It lets you think in terms of shared memory even when, as in the Compute-Server, no physically shared memory exists. Processes spun off by a Linda program can read, extract, or insert tuples, which resemble records in a database (see the text box "Meet Linda" on page 216 for a summary of Linda functionality).

For performance reasons, Torque implements tuple space as hash tables distributed among the ComputeServer processors. A client computer accesses tuple space through a "proxy" in the ComputeServer, which automatically converts all data types—including complex C structures—from the i860's format to the desktop computer's and vice versa.

A true computational server should fit transparently into desktop environments, not the other way around.

Programs running on your desktop computer and the ComputeServer look the same, but the ComputeServer can access tuple space faster—by orders of magnitude—than a desktop client can.

At the lowest level, ComputeServer processors access tuple space by using burst-mode transfers over the VT bus. The system divides tuples into small packets that it sends at speeds approaching the maximum transfer rate of the VT: 66 MBps. Round-robin arbitration gives every processor a turn to transmit a packet. Using packets to transmit tuples reduces latency by interleaving large and small transfers across the same channel. Otherwise, transferring large tuples (for example, a 10K-byte array) would generate unacceptable delays for other, smaller transfers waiting their turn.

To reduce hardware complexity, the transfer mechanism is actually implemented in software using the i860's register/memory block-copy instructions. The system first synchronizes the processors involved in a given packet transfer. Then the transmitting processor begins a block memory-to-register copy (thus reading from its memory). The synchronized receiving processors begin register-to-memory copies (writing to their memories) at the same time. As each processor increments its respective address bus, 64-bit-wide buffers simply "pour" the resulting data from one memory to another across one or more layers of the VT bus. Because the i860 is capable of driving even static column RAM at its full bandwidth, this method is as fast as any custom-designed DMA controller.

Optimizing Linda

A bare-bones Linda system is similar to a simple distributed database, with each processor in the system storing the tuples created by out() statements within its resident processes. When a request for a tuple can't be satisfied locally, the processor broadcasts the request to all the other processors.

If the request is an inp() or an rdp(), each processor replies immediately with either a matching tuple or a "no match" message. If an immediate match is not found, the remote processors store the request. The originating processor can periodically flush its old requests from the other processors. Aside from synchronization details, that's all there is to a basic Linda system.

One optimization employed on the ComputeServer is used by most Linda systems. It involves separating tuples into different partitions. In a given Linda program, out() operations put tuples into tuple space, while in() and rd() retrieve them from it.

Cursory inspection usually reveals that some out() statements produce tuples that can't possibly be retrieved by certain in() statements. For instance, out("Torque", 5.5) produces a tuple that in("Computer", ? & real) just won't read. Most Linda programs tend to have at least three or four natural partitions in their tuple space. Storing these separately improves the speed of the Linda run-time kernel on each processor.

A more sophisticated optimization is the pre-send. This technique uses the VT's hardware broadcast to send newly generated tuples to other processors before they are requested. No matter which process requests the tuple, it will find the information already resident on its processor. The principle is similar to that used by some disk-caching systems, which prefetch data blocks from the disk. In most cases, pre-send is best used sparingly, since the extra tuple copies it generates can rapidly fill system RAM.

The ComputeServer system decides which tuples to pre-send based on a mix of compile-time and run-time factors. For instance, it's better to pre-send tuples that rd() operations are going to match, since they leave the tuple around for other operations. The in() operation removes it.

The ComputeServer also minimizes network communications by having one of its processes act as a proxy for the desktop computer. All tuple searching and synchronization are performed between the proxy and the other processors in the ComputeServer. When a tuple is chosen, it's sent back to the desktop computer.

These kinds of optimization are critical in sustaining a sense of tuple space as a global resource, while at the same time providing you with a responsive system. If application code must be optimized around a system's shortcomings, then that system is meeting the Linda model in name only. In fact, an ideal multiprocessor system should go out of its way to accommodate "mistakes."

One example of this approach is dynamic load balancing. Because a distrib-

You can create client-server applications without buying ComputeServer hardware.

---

uted Linda kernel can both allocate time slices and monitor processes, it could actively seek out execution bottlenecks and take appropriate actions. For instance, an iterative process that repeatedly holds up others could be given successively larger time slices to minimize overall execution time.

Tools of the Trade
You can create client-server applications without having to buy ComputeServer hardware. The Torque developer's toolkit is a self-contained C-Linda imple-

mentation running on the client machine that allows code to be compiled, tested, and debugged as multiple threads on that system. The Linda primitives incur a very small speed penalty—typically 2 percent to 5 percent—when the application runs completely on the client.
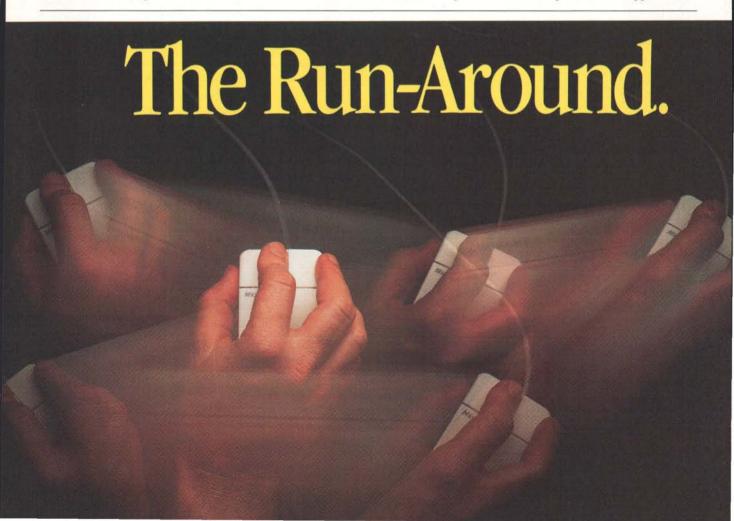
Thus, it's reasonable to maintain only one version of an application's source code that can be compiled for client-only or client-server operation. Once the code has been debugged on the desktop, it can be recompiled on the ComputeServer for client-server operation. Likewise, source code developed for the earlier transputer-based ComputeServer can be recompiled for the new machine without modification.

Mechanically, compilation for client-server operation is a matter of adding a few lines to the make file, because the ComputeServer compilers are themselves Linda applications that run from the desktop. Error messages from the Torque compilers are displayed within the standard programmer environments (MPW on the Mac, Microsoft's on the PC, and SunTools on the Sun).

The ComputeServer supports lan-

The Run-Around.

guages other than C-Linda. It does so in two ways: either by using a traditional, single-threaded syntax callable from a C-Linda framework, or as intrinsically parallel superstructures running above the Linda run-time environment. Examples of the former include FORTRAN 77 with VAX extensions and ISO Pascal. Examples of the latter are Paralogic's n/Prolog, which is a multiprocessor interpreter written in C-Linda, and Strand's Strand88, a parallel language with sophisticated job-control functions. Finally, several high-level graphics tools are under development for the system, including the PPSE parallel CASE tools developed at Oregon's Advanced Computing Institute.

Join the Party
Torque has concentrated its efforts on applications that combine an advanced user interface with sustained number crunching. Most of these involve 3-D graphics, simulation, and image processing. Several prominent vendors have pledged support for the system, including Macro-Mind (Three-D 1.1) and Wolfram Research (Mathematica kernel). Smaller

niche vendors are also supporting the system, including Market Engineering (Crystal Ball, for Monte Carlo forecasting), and Pre-Press Technologies (SpectreSets color-separation software).

ComputeServer-capable versions normally cost more than their desktop counterparts, but they provide considerably higher performance and support multiple users. A number of embedded applications are also in the works, including exposure control for submicron lithography at Lepton (an AT&T spin-off).
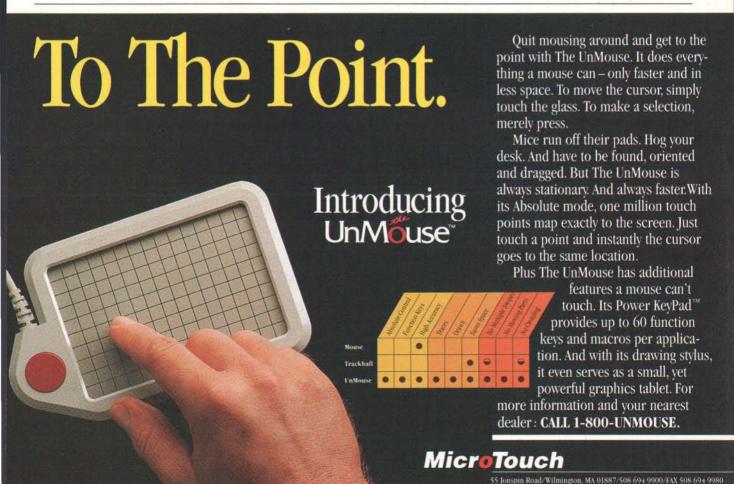
Third-party ComputeServer developers represent a collection of seemingly unrelated specialties, with little in common except the need for speed. In fact, support for this system by "mainstream" applications is almost an oxymoron: If a product performs well enough to be sold by the millions on today's desktops, it probably does not need the Compute-Server. Rather, we expect that the ComputeServer will help bring today's niche products into tomorrow's mainstream.

There's nothing to keep ad agencies from ray-tracing animated sequences if each frame takes only minutes to produce. There's nothing to keep financial

analysts from running 100 Monte Carlo variations on a spreadsheet if the job takes only slightly longer than a single recalculation. And there's nothing to keep desktop publishers from running color separations alongside Linotronic output when the job no longer takes hours to complete. Such applications will become commonplace when access to the necessary computing power becomes available.

A new kind of computer, the computational server, will complement desktop machines for compute-intensive applications. By treating computational power as a shared resource, the ComputeServer delivers lots of FLOPS at a reasonable price per desktop, and it does so without forcing you to sacrifice the computing environment you're comfortable with. Given the ever-increasing importance of LANs, the ComputeServer is truly a machine for the 1990s. ∎

Sam Bogoch, Iain Bason, Jeff Williams, and Mike Russell design and develop computational servers for Torque Computer, Inc. (New York). They can be reached on BIX as "sbogoch."

New C 6.0. It's not

for tachophobiacs.

If you love the rush of adrenalin that comes with sudden acceleration, we've got a C that's just your speed: the exhilarating new Microsoft® C 6.0 Professional Development System.

Designed to get your blood racing in nothing flat with the fastest,* slickest code available in the shortest, smoothest time possible.

Thanks to some of the most sophisticated code optimizations around.

With this C, you get everything from register-based parameter passing to a globally optimizing code generator to complete loop optimizations. Plus new super-efficient based pointers that access far data instantly.

And your code isn't all that moves faster.

Whether you're developing for MS-DOS,® Microsoft Windows™ or MS® OS/2 Presentation Manager systems, our brilliantly integrated Programmer's WorkBench propels you through the creative process with precision tools.

Including a new souped-up CodeView® Debugger that not only debugs any size DOS or OS/2 application program, on any 286 or 386 machine, but also gives you data-browsing so you don't have to constantly guess at values.

To save even more time, our C Advisor is on-line and on-time whenever you need it. This complete hypertext reference not only gives you sample coding solutions, but it even lets you copy and paste them directly into your program.

If you want to review part of your handiwork, our Source Browser helps you find it with a full call tree that literally draws you a map.

For a free white paper with more details, just call us toll-free at (800) 426-9400.

Then get your hands on Microsoft C 6.0.

The development system specially created for people whose only fear is being left behind.

Microsoft C Professional Development System

Compiler Performance:
- New based pointers access far data with size and speed of a 16-bit pointer.
- Register parameter passing: _fastcall passes up to 3 parameters in registers, and floating point values are returned on the coprocessor stack.
- Function-level control over optimizations.
- Integrated inline assembler.
- Highly compliant with proposed ANSI standard (we pass 75 of 79 "Plum Hall" tests).
- Incremental compilation in both DOS and OS/2.
- Dramatically improved local code-generation.

Programmer's WorkBench:
- Designed by professional programmers, for professional programmers.
- Complete integration of edit, make, debug cycle in both DOS and OS/2.
- New Source Browser provides information on all aspects of your source code.
- Microsoft C Advisor is an on-line reference for the environment, C Language, and Runtime Libraries.
  - New MAKE is now a full superset of XENIX/UNIX MAKE facility.
  - Specify all build options, including target environment, from within development environment.
  - Takes advantage of OS/2 for background operation of components—such as compilation.
  - Open architecture: 3rd party tools and libraries can be easily installed.

Microsoft CodeView Debugger:
- Completely redesigned user interface.
- Debugs nearly any size application on both 286 and 386 machines under DOS by taking all but 15K of CodeView out of the 640K address space.
- Multiple file debugging.
- Multiple memory views.
- Array browsing.
- Automatic "locals" window.
Advanced Programming Techniques Manual answers the most popular programming questions with extensive commentary and sample code.

Microsoft®
Making it all make sense™

In the 1990s, code will be generated by the click of a mouse or a tap of a key. With Matrix Layout 2.0 you can do that now. And the results will surprise you.

Preview the 1990s with Layout

In Layout, you create programs by designing an object-oriented flowchart, with all the options of traditional programming. It's a technology we call desktop programming.

Once you're done, simply choose the language you want for the finished program. There's Microsoft C, Lattice C, and Turbo C, as well as Turbo Pascal and Microsoft QuickBasic. You can even create a .EXE file that's ready-to-run on any IBM PC or compatible.

1990s Power without 1980s Pain

Because Layout works with today's standards, you can painlessly take advantage of the power behind Layout – object oriented programming, CASE (Computer Aided Software Engineering) technology, hypertext databases, and graphical user interfaces. All without giving up your favorite computer language.

An Architecture for the 1990s

Layout comes with objects that produce real code for everything traditional computer languages can do – math, branching, variable management, complex data structures – and it extends each language to include powerful user interface and hypertext database capa-

bilities. But best of all, you can extend Layout past the 1990s by building your own objects – BlackBoxes – that can do anything you imagine. Added together, Layout cuts your development time by up to 70%.

Welcome to the 1990s
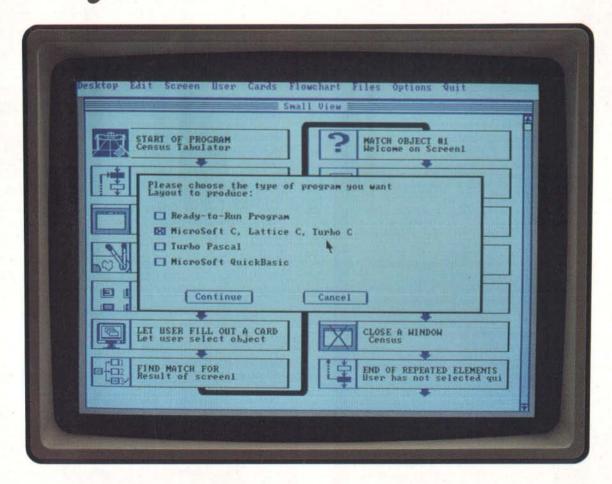
Ready to jump into the 90s? Get Layout today. It's available for just $199.95. Or for a glimpse of the 90s, see the Layout video tape for just $9.95. Give us a call at

1-800-533-5644

In Massachusetts call (617) 567-0037.

MATRIX SOFTWARE

In the 90s, this is how you'll write code.

Join the EISA Evolution

*The EISA bus is breaking up that
old "Gang of Nine"*

Min-Hur Whang and Joe Kua

S upercomputing has a definite ring to it. It even has a magical quality. It makes you think of the HAL 9000 or of buildings full of giant computers that make the earth spin or change gravity. But when you analyze the concept of supercomputing, you find a combination of elements, each of which contributes to the whole.

The goal of any supercomputing architecture is to obtain increased performance through increased system throughput. One of the more important ways to increase throughput is to free up your processor (or processors) from such mundane tasks as system I/O. Supercomputers use either a separate I/O processor or a separate computer (often a minicomputer) to handle I/O to and from main memory.

Until the advent of sophisticated buses such as the Micro Channel architecture from IBM, Extended Industry Standard Architecture from a collaborative effort nicknamed the "Gang of Nine" (AST, Compaq, Epson, Hewlett-Packard, NEC, Olivetti, Tandy, Wyse, and Zenith), and NuBus from Apple, personal computers handled I/O the old-fashioned

way, by moving every byte through the processor when performing I/O with main memory. Now, however, the new buses perform I/O without tying up the CPU. The EISA bus master is the latest example of using such a scheme to improve I/O for microcomputers.

EISA, as the name suggests, is an extension of the original Industry Standard Architecture bus (the PC AT bus). Since

EISA evolved from ISA, it kept all the preceding technology. Thus, you can use your current IBM-compatible expansion cards on an EISA machine. The Micro Channel architecture doesn't offer this evolutionary approach, so simple items like network cards, SCSI adapter cards, and modems have to be specially designed to support its I/O bus. EISA, however, is already supported.

A Backward Glance

The system bus for the original IBM PC was an 8-bit extension of the Intel 8088 microprocessor. The 8088 has only 20 address lines, restricting the amount of addressable memory on the PC to about 1 megabyte. Since today's memory-hungry operating systems (e.g., OS/2, Unix, and Novell NetWare) require more memory than that, the 8-bit system has become obsolete. In addition, the original 8-bit system bus offered only an eight-level edge-triggered interrupt and four DMA channels whose transfer speeds ranged from about 250,000 bps to 500,000 bps.

The next generation of bus was the 16-bit PC AT bus, which became the ISA

continued

Requesting A Transfer

An EISA bus master supports two types of data transfer cycles for moving data between itself and main memory or an I/O slave. The cycles are either standard or burst. A third cycle type, the compressed cycle, is available only to the CPU. EISA cycle types are summarized in table A.

In a standard cycle, an EISA bus master is completely synchronous with the bus clock (BCLK), which provides synchronization with the main system clock. The bus clock usually operates at frequencies of between 8.33 MHz and 6 MHz, although its period is sometimes extended to synchronize with the host CPU or other motherboard devices.

An EISA bus master drives nine signal lines on the bus and monitors five others (see table B). Using these signals, you can examine how a bus master transfers data to memory or to I/O slaves within a standard cycle. To follow the events in a standard transfer cycle, the numbers in the discussion below refer to the circled numbers in the timing diagram in figure A.

Let the Games Begin

When a bus master wants to initiate a transfer, it requests control of the bus by asserting the $\overline{\text{MREQx}}$ signal (1). It then monitors the $\overline{\text{MAKx}}$ line (2) until the EISA arbitration system grants its request (2').

Next, the bus master places the target address on the bus and indicates whether it is initiating a transfer with system memory or with an I/O slave (3). Then, it indicates whether it is reading from or writing to the indicated address, and which bytes of the 32-bit-wide data path it wants to transfer (4).

At the same time, the master drives the $\overline{\text{START}}$ signal to initiate the transfer cycle (5). At the end of the $\overline{\text{START}}$ signal, which lasts one BCLK cycle, the bus master negates the $\overline{\text{CMD}}$ line (6).

CYCLE TYPES AND TIMES

Table A: *EISA supports three types of bus transfers, of which only two—standard and burst—are available to bus masters. This table lists the average number of bus clock cycles it takes to complete a data transfer cycle. (N/A = not applicable.)*

Cycle	Speed	Standard cycle	Compressed cycle	Burst cycle
EISA	32-bit	2	1.5	1
One-wait-state ISA	16-bit	3	N/A	N/A
	8-bit	6		
Zero-wait-state ISA	16-bit	2	N/A	N/A
	8-bit	3		

EISA STANDARD TRANSFER CYCLE
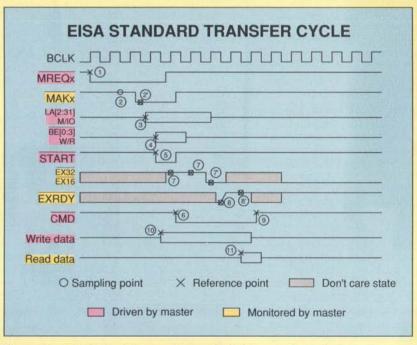
Figure A: *This timing diagram of a standard data transfer cycle shows how the master and slave use signals to synchronize their activity. When writing, the master dumps the data onto the bus when it starts the transfer. The slave reads the data and indicates when it is through using the $\overline{\text{EXRDY}}$ line. When reading, the master waits until the slave indicates it is done before sampling the data bus.*

bus. Improvements in the microprocessor (with Intel's 286) and in the number of addressable memory locations (with 24 address lines) brought the total available memory to 16 MB. The edge-triggered interrupts increased from eight to 15 levels, and the DMA channels from four to seven.

The speed of the 286 supports 8-bit and 16-bit modes with DMA data trans-fer rates of 100,000 bps to 400,000 bps, and 400,000 bps to 800,000 bps, respectively. Note that the new 8-bit DMA mode is slower than the original PC bus. Intel had to sacrifice some speed to maintain downward compatibility to 8-bit devices. One benefit of the 16-bit bus is that you can continue to use any 8-bit cards you have, because ISA is a superset of the 8-bit PC bus.

From a technological point of view, the problem of I/O bus design became obvious with the arrival of the Intel 386 and i486 processors. ISA just doesn't cut it. An I/O bus for these applications workhorses needs the following additional capabilities:

• more memory capacity;
• multiple bus-master devices with

EISA BUS-MASTERING SIGNALS

Table B: *The EISA bus features special signals that enable a master to take control of the bus. If both the master and slave are burst-capable, the signals also enable burst transfers.*

Controlled by a master

$\overline{\text{MREQx}}$	Asserted by a bus master when requesting use of the bus. The x is the master's slot number.
$\overline{\text{MAKx}}$	Indicates that the master in slot x has control of the bus.
LA[2:31]	Address lines.
M/IO	Indicates whether the transfer involves main memory or a slave device.
$\overline{\text{BE[0:3]}}$	Indicates which bytes of the 32-bit-wide bus are used in the current bus cycle.
W/R	Indicates whether the master is writing or reading data.
$\overline{\text{START}}$	Indicates the start of a bus-mastering cycle.
$\overline{\text{CMD}}$	Asserted by a master during a transfer cycle.
$\overline{\text{MSBURST}}$	Indicates that a master is capable of controlling burst transfers.

Monitored by a master

D[0:31]	Data lines.
$\overline{\text{EX32}}$	Indicates that the slave handles 32-bit transfers.
$\overline{\text{EX16}}$	Indicates that the slave handles 16-bit transfers.
$\overline{\text{EXRDY}}$	Indicates that the slave has completed its function and is ready to terminate the cycle.
$\overline{\text{SLBURST}}$	Indicates that a slave is capable of burst transfers.

Following the confirmation of $\overline{\text{START}}$, the bus master samples $\overline{\text{EX32}}$ and $\overline{\text{EX16}}$ (7) to determine whether the slave is a 32-bit or a 16-bit device. Once it determines the type of slave (7'), it checks for $\overline{\text{EXRDY}}$ from the slave. When the slave is ready to terminate a transfer cycle, it drives the $\overline{\text{EXRDY}}$ signal (8).

When the master confirms $\overline{\text{EXRDY}}$, it negates $\overline{\text{CMD}}$ to finish the cycle (9).

When writing, the data from the bus master is valid once it drives the $\overline{\text{START}}$ signal (10). When reading, the master samples the data lines on the first rising edge of BCLK after the slave asserts

continued

high-speed burst-transfer rates;
- a 32-bit data transfer rate for the CPU, DMA, and bus-master devices;
- enhanced DMA arbitration and transfer rates;
- level-triggered (shareable) interrupts; and
- automatic configuration of system and expansion boards.

To promote these features, two completely new bus architectures, IBM's Micro Channel and EISA, were introduced. A generation gap exists between EISA and the current reigning ISA buses.

The Bus Stops Here
When the specifications for implementing EISA were released, the bus standard attracted a lot of attention. The intent of

EISA was to develop a 32-bit architecture that, unlike the Micro Channel, would be downward-compatible with the existing AT bus. Systems using EISA would still be able to use 8-bit and 16-bit boards. The EISA standard adds many functional and performance enhancements, such as improvements in DMA handling, but it retains a subset of the features available in the ISA bus.

EISA's data path can be 8, 16, or 32 bits wide. The number of address lines is increased to 32, which theoretically allows access to up to 4 gigabytes of memory, depending on the system design. This larger amount of addressable memory better meets the needs of today's memory-hungry applications, such as multiuser systems, network-server systems, AI CAD/CAM simulations, and even computer hardware design.

Not only can the CPU fully implement EISA's entire 32-bit address space, but so can the 32-bit bus-master and DMA devices on the EISA bus. EISA's 32-bit DMA functions allow today's 16-bit ISA DMA products to use the 32-bit address space by adding a simple software device driver. In other words, you can keep your system's old equipment *and* improve its performance.

The EISA bus master includes a dedicated I/O processor and local memory. This specialized processor drives the address, data, and control signals for intelligent peripherals, which become slave devices, during a bus cycle. Bus masters improve system performance by taking on simple tasks that would otherwise fall to the host processor. Thus, they reduce the main processor's work load (the beginning of multiprocessing, perhaps?). To understand the benefits of the bus master, it is important to grasp the mechanics of transferring data between the bus masters and their slave devices (see the text box "Requesting a Transfer" at left).

The capabilities of the EISA bus master are aimed at sophisticated I/O peripherals that require optimum performance or advanced memory-access functions (e.g., non-ordered scatter/gather data operations). As a result, EISA is primarily intended for 32-bit devices, which typically contain dedicated I/O processors and require the fastest data transfer rate available through the bus. For this purpose, a fast burst-transfer mode is available to bus masters, in addition to the more typical timing of I/O and memory cycles.

In one scenario, a system could have several bus-master peripherals running

continued

EXRDY (11).

If the slave does not reply immediately to the bus master, it will wait another BCLK cycle. In the example, EX16 or EX32 missed the call from the bus master, thus incurring two wait states (7), and the master missed EXRDY, adding another wait state (8'). The result is a three-wait-state cycle. A standard cycle with no wait states lasts only two BCLK cycles.

Burst of Power

In a burst cycle, a burst-capable bus master uses all the signals defined for the standard bus master, in addition to driving the MSBURST signal and monitoring the SLBURST signal from the slave device. The numbers in the discussion below refer to the circled numbers in figure B.

A burst cycle follows the same sequence as a standard cycle until after the master asserts START. At this point, the master begins sampling SLBURST, as well as EX32 and EX16. As before, once the slave affirms EX32

or EX16, the master starts sampling EXRDY. However, if the slave is burst-capable, it asserts SLBURST (1) in addition to EX32 or EX16. When this happens, the master negates MSBURST (2) and begins the burst cycle.

As before, data from the master is valid when START is asserted (3). Data from the slave is valid (4) until CMD ends the cycle.

After the cycle begins, the bus master provides the next LA[2:31], M/IO, BE[0:3], and W/R when it samples EXRDY. It also samples MSBURST at every rising edge (5) unless EXRDY is not accessible (6). For the burst cycle, each wait state must have one BCLK cycle.

The master ends the cycle by disasserting MSBURST (7) and inverting CMD (8). As with the standard cycle, CMD is first negated at the end of the START signal. In a zero-wait-state system, a burst-mode cycle will take only one BCLK cycle with the exception of the initial cycle, which requires two BCLK cycles.

on a network file server, each dedicated to a special application, such as distributed database processing. EISA can support up to 15 bus masters with an arbitration method that determines the latency of each device. Using this latency, the response time for requests from expansion-bus devices can be determined.

Knowing Your ABCs

In EISA systems, DMA devices have seven channels, just like they do in ISA systems, but the EISA transfer rate is much faster. The DMA controllers support 8-, 16-, and 32-bit data transfer sizes. They have four cycle-control sequences for transferring data between the DMA device and memory. These four cycles are

- the ISA-compatible cycles, which execute one transfer cycle in eight bus cycles;
- Type A cycles, which execute one transfer cycle in six bus cycles;
- Type B cycles, which execute one transfer cycle in four bus cycles; and
- Type C (burst DMA) cycles, which execute one transfer cycle in one bus cycle.
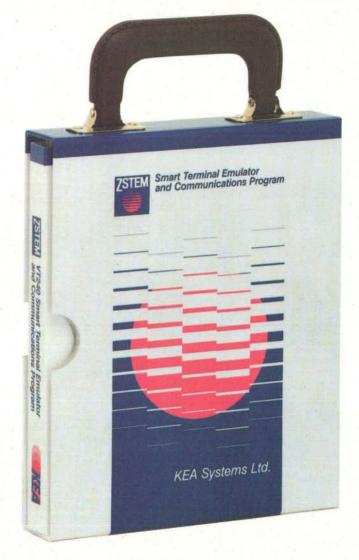
Types A, B, and C support 8-, 16-, and 32-bit DMA devices and perform automatic data-size translation for transfers to mismatched memory. Table 1 indicates peak data transfer rates and compatible DMA devices for each DMA cycle type. Moreover, most ISA-compatible DMA devices can transfer data 130 percent to 200 percent faster by programming the EISA controller to use Type A and B transfers instead of ISA-compatible timing.

How is this increased performance and efficiency provided? By enhanced arbitration, which shortens the time between the DMA device's request and grant events. This enhancement does not imply a decrease in compatibility. Existing hardware and software can take advantage of it without modification, so it actually improves compatibility with older systems.

Fringe Benefits

The original PC and ISA buses used edge-triggered interrupts, which are easy to implement but susceptible to false triggering and cannot be shared with other peripherals. In addition to supporting these edge-triggered interrupts to maintain compatibility, EISA also provides level-triggered interrupts, which are less susceptible to noise and allow

EISA BURST MODE

O Sampling point   X Reference point   ▢ Don't care state

▢ Driven by master   ▢ Monitored by master

Figure B: *Whereas standard mode requires that the master and slave synchronize for every 16- or 32-bit transfer, burst mode requires that they synchronize just once, before the first transfer. The average cycles per transfer in burst mode is thus very close to one.*

Our VT240 terminal emulator has changed the meaning of portability.

What does VT240 portability mean to you?

- **Portable across operating systems!** Our ZSTEM 240 software runs under both MS-DOS and Unix 386 System V, making your PC look and act like a VT240/340 terminal on either operating system.

- **Portable across machines!** ZSTEM 240 runs on IBM PCs, XTs, ATs, PS/2s and compatibles, from AT&Ts to Zeniths!

- **Portable across portables!** ZSTEM 240 runs on portable PCs, letting you connect to your office systems when you're on the road.

- **Portable across video adapters!** ZSTEM 240 supports all standard video adapters: VGA, EGA, CGA, MCGA, AT&T, Hercules and many extended adapters. No matter what adapter/monitor combination you use, ZSTEM 240 displays double-high/double-wide characters, 132 columns, and VT340 ReGIS, sixel and Tektronix graphics.

- **Portable across networks!** ZSTEM 240 connects to your favorite networks, including Novell, 3COM, TCP/IP, Ungermann-Bass, Excelan, Wollongong, FTP, Sun and DEC's CTERM and LAT.

Of course, ZSTEM doesn't *really* come with a handle. What it does come with is our top-notch technical support and documentation, plus a solid warranty, so you can be assured of quality products backed by quality people. Call today about our complete line of VT emulation products.

KEA Systems Ltd.
3738 North Fraser Way, Unit 101
Burnaby, B.C., Canada V5J 5G1
Telephone: (604) 431-0727
FAX: (604) 431-0818

Toll-Free Order Desk
800-876-6089

See us at DEXPO South Booth #967

EISA also includes a product-identification mechanism for systems and expansion-board products. As you turn the system on, the computer automatically runs a power-up sequence of self-tests. When it does, it interrogates each device for the product identifier, compares that with those stored in RAM, and configures the board according to the setup data stored in ROM.

Micro Channel vs. EISA

The Micro Channel's design is similar to that of EISA, and it includes support for the 386 and i486 microprocessors. It also supports up to 15 bus masters and can transfer data in burst mode. The Micro Channel also replaces DIP switches and jumpers with a self-configuring system. Unlike the ISA bus, the Micro Channel's level-triggered interrupts and sophisticated DMA-arbitration scheme allow multimaster operation. The major disadvantage of the Micro Channel, however, is that it is not downward-compatible with the ISA bus.

Supporters of the Micro Channel architecture argue that you won't want to migrate your original expansion cards to new machines. However, even if you don't, there are tens of thousands of application programs and thousands of expansion boards on the market today that EISA supports that the Micro Channel does not. In addition, most Micro Channel–based systems have basic features built into the machine, such as graphics adapters, mouse/serial/parallel ports, and floppy and hard disk controllers. These built-in features limit your future choice of products.

Micro Channel systems transfer data between the CPU, memory, and 32-bit peripherals at a maximum rate of 20 megabytes per second. That's less than two-thirds the speed of an EISA system, which normally performs at about 33 MBps. The Micro Channel offers a maximum of eight expansion slots, compared to EISA's 15.

Physically, the form factor of a Micro Channel board is about half the size of an EISA board. This makes Micro Channel product design more difficult and expensive. Since the board is smaller, the circuitry must also be smaller. In some cases, this means expensive application-specific ICs or surface-mounted ICs are necessary.

In addition, the Micro Channel's card has less than half the power of an EISA card available. This makes peripherals, like I/O boards with large amounts of memory, much more complex and ex-
continued

multiple peripherals to share the same interrupt level. Theoretically, level-triggered interrupts can have an infinite number of levels.

Another practical benefit that EISA provides is automatic configuration of system resources and expansion boards.

That means an end to DIP switches, jumpers, and installing configuration files. A system that supports plug-and-play peripherals improves efficiency (switch configurations for ISA or EISA products are still allowed, if you need them).

DMA CYCLE TYPES

Table 1: *Most ISA-compatible DMA devices can transfer data faster if the EISA controller is set up to use Type A and Type B transfers instead of ISA-compatible timing. Enhanced arbitration shortens the time between the DMA device's request and grant events.*

DMA cycle type	Size of transfer	Transfer rate (MBps)	Compatibility
Compatible	8-bit	1.0	All ISA
	16-bit	2.0	All ISA
Type A	8-bit	1.3	Mostly ISA
	16-bit	2.6	Mostly ISA
	32-bit	5.3	EISA only
Type B	8-bit	2.0	Some ISA
	16-bit	4.0	Some ISA
	32-bit	8.0	EISA only
Burst DMA (Type C)	8-bit	8.2	EISA only
	16-bit	16.5	EISA only
	32-bit	33.0	EISA only

pensive to design and build for Micro Channel systems.

EISA Takes Its Stand

Intel has already introduced EISA chip sets, including three chips for the motherboard and one chip for an add-in card. The first motherboard chip has been designated the integrated system peripheral and provides 32-bit DMA control, timer/counter control, interrupt control, bus arbitration, and DRAM refresh functions on a single chip.

The EISA bus controller, the second chip in the set, interfaces with the 8- and 16-bit ISA bus, the 32-bit EISA bus, and the CPU by performing the timing and control functions. The last motherboard chip, which is called the EISA bus buffer, contains buffering logic for any one of the data, address, and parity-control modes.
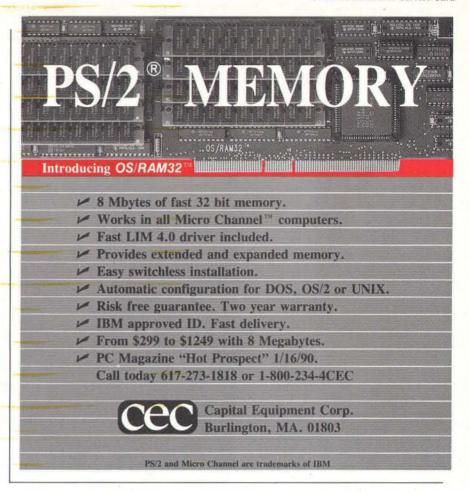
The intelligent add-on cards chip, which is called the bus-mastering interface controller, serves as the interface between a plug-in card and the 32-bit EISA bus and provides interface logic, drivers, and a DMA controller that is capable of becoming a bus master. EISA chip sets work with the 386 and i486 CPUs at speeds of up to 33 MBps.

The members of the original Gang of Nine are starting to introduce EISA systems, but not all of them are ready yet. Only a few currently have EISA systems. Hewlett-Packard was the first to produce an EISA microcomputer system, the Vectra 486PC. It transfers data at a rate of 20 MBps with a 16-bit ISA ESDI controller.

Right before Fall Comdex 1989, Compaq revealed its EISA system, the Systempro. The Systempro can decrease the time required to transfer a file to or from the disk by distributing a single file among multiple bus-master controllers, called an Intelligent Drive Array (IDA). Compaq also claims that its system will transfer data at 33 MBps.

Zenith Data Systems' 386/33E is the latest EISA PC system. The bus-master controller, which uses 1 MB of memory and is expandable to 4 MB, will support up to four ESDI devices, seven SCSI devices, and two floppy disk drives. The Zenith 386/33E is also supposed to be able to support a data transfer rate of about 33 MBps.

Some companies outside the original Gang of Nine are also working on EISA products. By the end of June, Arche Technologies will expand its current 486 AT system (up to 64 MB of memory, 256K bytes of external cache, and enhancements for write posting, page

memory cycle, and prefetch cycles) to support the EISA standard. The Arche 486 EISA system will be able to use 32-bit VGA, 32-bit Ethernet, and 32-bit SCSI host adapters.

Although EISA systems are downward-compatible with ISA expansion boards, a number of EISA-specific products are expected to appear by the end of this year. These include

- Adaptec's 32-bit SCSI controller,
- Distributed Processing Technology's 32-bit hard disk controller,
- Proteon's 32-bit EISA version of a token-ring adapter,
- Standard Microsystem's 32-bit EISA version of an Ethernet board,
- SunRiver's 32-bit fiber-optic board,
- 3Com's network card,
- Arcnet's 16-bit multiuser board, and
- Western Digital's 16-bit ESDI controller.

EISA systems will thrive on the availability of EISA-specific products and the already existing stock of ISA products from current PC machines.

Make New Friends but Keep the Old

EISA will save you money, time, and paperwork because it will allow you to

make better use of the hardware and software resources you already have. Like its predecessor, ISA, EISA is an evolutionary approach to the need to improve system performance. It offloads I/O processing from the microprocessor's list of things to do.
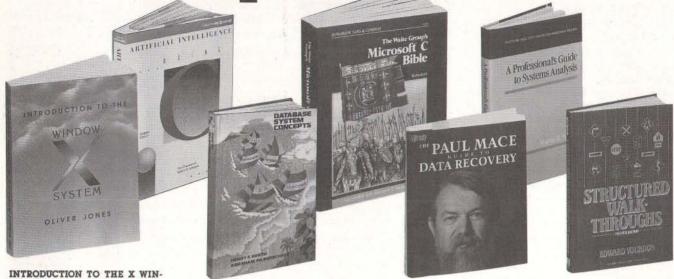
Unlike the Micro Channel, however, EISA does not require new peripheral cards to support its I/O bus. It uses standard AT cards. While adapting to the current AT bus architecture and using the more powerful 32-bit processors, EISA preserves compatibility with the existing world of AT-compatible devices. EISA has it all. ∎

Min-Hur (Michael) Whang is the technical specialist at Arche Technologies, Inc. (Fremont, CA). He has a Ph.D in electrical engineering from Polytechnic University of New York. Joe Kua is the director of research and development at Arche. He was the chief designer of Everex's Step machines. They can be reached on BIX c/o "editors."

A Calculating RISC

Coprocessors based on RISC engines will soon deliver supercomputing performance to your desktop

Trevor Marshall

Many personal computers can already deliver the same computational performance as mainframes, if you add in their coprocessing capabilities. With the advent of RISC coprocessors, these systems will approach speeds that you now associate with supercomputers. This next leap in performance will bring some unique advantages to the coprocessing environment.

Who Needs It?

When you think of a supercomputer user, you probably conjure up an image of a scientist in a research laboratory laboring over a massive problem in nuclear physics. While it's true that this type of user tends to have a virtually insatiable desire for compute power, the bulk of high-performance computing is devoted to three-dimensional graphics, mathematical simulations (e.g., aircraft flight simulators), and routine data analysis (e.g., in geophysics). Even apparently mundane applications in the printing industry, such as optical character recognition, color prepress, and typographic output, make up a significant sector of high-performance computing.

The applications that benefit most from coprocessing technology are those that are "compute-intensive"—where most of the processing is actual number crunching or manipulating data in memory. Because the host computer is used for disk I/O and operator interaction, a coprocessor is not very effective in situations where its CPU spends a significant amount of its time waiting for I/O tasks to complete. Thus, applications involving considerable file manipulation, for example, will benefit least from using a coprocessor.

Host Machines

Coprocessors are generally designed to boost the performance of workstations and high-end personal computers. Workstations are usually the most expensive desktop computers, and they often feature a multiuser capability. High-end personal computers are typically DOS-based systems using 386 or i486 CPUs with 80387 or Weitek floating-point chips. The Mac II family has also become a popular high-end platform, due to its intuitive user interface.

Both workstations and personal computers are mass-produced as general-purpose machines. Coprocessors are complete computing systems that plug into the expansion bus on these machines and work with their 386, 68020, or RISC CPUs, offloading those tasks that the general-purpose machines don't do well. This allows the host to concentrate on the I/O, file system, and operator interface while the coprocessor does the number crunching.

continued

Photo 1: *This frame from the He-Man animation, created by Mr. Film using software from Digital Arts, shows how the power of a coprocessing system translates into tangible graphics. (Image copyright 1989 by Mattel Corp.; used by permission.)*

Coprocessing Through the Ages

Like desktop computers, high-performance coprocessors are a relatively recent phenomenon. The first was Steve Ciarcia's Trump Card, described in the May and June 1984 BYTE. The Trump Card featured a Zilog Z8000 16-bit CPU and plugged into an IBM PC. Although the available software was limited to a BASIC interpreter, the Trump Card nonetheless defined many of the features of today's high-powered coprocessing systems.

In the August and September 1985 BYTE, I described the Definicon Systems 32032-based coprocessor, which improved on the Trump Card in two ways: It used a true 32-bit microprocessor and came with a fully defined MS-DOS software interface. With multiple languages (C, Pascal, and FORTRAN) available that had been ported from the 32000 Unix environment, the Definicon DSI-32 coprocessor offered a high degree of software portability between Unix and DOS systems. The interaction between the compilers and the DOS operating system had evolved to where the coprocessor could execute even 80x86 software interrupts, and thus access all the features of the PC.

The following year, a 68020-based coprocessor system (the DSI-020) finally achieved a performance level exceeding that of minicomputers. This system ensured the success of the coprocessor concept, and it still forms the basis of software systems as complex as nuclear-reactor simulations.

A Stellar Performance

Coprocessors have several advantages over their workstation and personal-computer hosts that explain the performance differential between them. First, the CPU on the coprocessor usually doesn't run a complete operating system (like DOS or Mac OS), but operates in a shell of the host operating system.

You usually invoke an application running on a coprocessor with commands or mouse activity, the same as you would to execute a native application. Consequently, when a new CPU is released, coprocessor companies need only rewrite the shell code to take advantage of it—a task that usually takes just a few months.

Thus, new CPUs normally appear in coprocessor form as much as a year before they appear in a complete system implementation. Given the rate of performance increases in microprocessors over the last decade, this edge frequently translates into a 300 percent advantage in performance for coprocessors versus complete systems at any particular point in time.

Also, you can usually mount multiple coprocessors on one host and give each a different task to perform. This is especially useful with 3-D graphics software.

Finally, a coprocessor rarely has more than a rudimentary expansion bus. Unlike a conventional desktop system, all the memory and I/O subsystems on a coprocessor are closely coupled and optimized for better performance rather than compatibility with an existing expansion structure.

Boards of Distinction

Many people make no distinction between array processors and true coprocessors. Array processors, however, are generally designed for specialized tasks, such as computing fast Fourier transforms (FFTs), and come with libraries of subroutines that are linked with the main code running on the host. Most programs have to be rewritten to take advantage of an array processor.

A computing coprocessor, on the other hand, has extensive software support (including compilers) to minimize or eliminate the need to rewrite applications software. Generally, all you need to do to complete a port to a coprocessor is recompile the source code, using the software tools supplied with the coprocessor.

Array processors have advantages in two primary areas. First, if the computing task is quite small but repetitive (such as an FFT), an array processor can usually compute it faster than a general-purpose CPU.

Second, where a large proportion of the application is interaction with the operator via a complex operating environment (as in the Macintosh), it may be easier to let much of the code run untouched on the host and link in a few key subroutines running on the array processor. Several coprocessor manufacturers have designed linkable interfaces to their boards to allow them to be used this way as well.

The Soft Sell

Little third-party software is available for most coprocessors. One exception is transputer-based systems, which have a proliferating base of applications software that generally runs on any PC-based transputer board. Most of the available third-party software comes from specialized developers whose applications require a lot of computing power. Otherwise, much coprocessor software development remains in the hands of end users.

If you're writing coprocessor software rather than using third-party products, the software supplied to compile and run your application will be the most important part of your system. Be aware that the quality of compiler support varies

greatly from system to system and even between identical CPUs on different platforms.

Although software for coprocessors is not nearly as abundant as it is for desktop computers, the programs available perform tasks beyond the scope of standard personal computers and workstations.

One natural application for coprocessing is graphics rendering. Recognize the image from the He-Man animation in photo 1? It was computed with coprocessors. In fact, coprocessors compute the animation for many of the title sequences you see on TV.

The DGS system that was used to create the He-Man image was created by Digital Arts. During product design, Digital Arts quickly realized that the PC itself did not have the power or the architecture to run complex rendering code at commercially viable frame rates. Each image in such an animation sequence takes about 45 minutes to compute, even on a coprocessor-equipped system.

The total computing time needed to produce the 2-minute video (at 15 frames

Photo 2: *Although written for 680x0 machines, the Oklahoma Seismic MIRA program runs on DOS machines, thanks to coprocessor technology. The program gives geologists and petroleum engineers a view of the structural geology of a region.*

The fact is, coprocessors compute the animation for many of the title sequences you see on TV.

per second) was about eight weeks. If you consider that the DGS system uses two coprocessors—a 68030 to set up the images and edit the sequence, and a transputer to perform the actual rendering—you can appreciate that graphics rendering is computationally intensive indeed. Without coprocessors, the time required would escalate dramatically.

The graphics card and video interface that the DGS system uses are standard PC devices. The DGS software allows the results of computation on the coprocessor to be sent to the video interface.

Another point is worth noting here. There is a certain amount of standardization in coprocessor technology. Coprocessors not specialized to perform array processing or graphics rendering are designed to take standard source code (usually from a Unix system) and then com-

pile it with minimal changes to run on a personal computer. Therefore, once you have started using coprocessor technology, it becomes a relatively straightforward task to switch to a new processor when it becomes available. This tends to keep competition strong and prices reasonable.

The DGS software, for example, uses a 68030 coprocessor from either Definicon or YARC Systems, and transputer cards from CSA, Microway, or YARC.

Powerful Pluses
There are many reasons to use coprocessors besides raw power. One is to ensure object code compatibility. A coprocessor is probably the only way to market an application that uses, for example, 68020 assembly language in the Intel-dominated PC world, short of a total rewrite of the code.

This is why Oklahoma Seismic used a 68020 coprocessor for its MIRA geophysical data-analysis programs. MIRA analyzes the echo returns from sound blasts set off at regular intervals along the earth's surface that estimate a probable reconstruction of the underground geophysical-layer structures. A typical reconstruction is shown in photo 2. No compromises need to be made in a program's graphical or operator interface because it is running on a coprocessor.

Another attraction of coprocessors is the ability to increment your computing power just by adding more boards. For example, Visual Information (VIDI) produces an animation package for the Mac II that is similar to the Digital Arts

DGS software. The VIDI software running on a single 29000 RISC coprocessor runs about 10 times faster than on a standard Mac II, and about 5 times faster than on the Mac IIci.

VIDI has written some software that allows up to five RISC boards to simultaneously render the images from an animation sequence. This gives an effective system speed 50 times greater than that of the Mac II host itself. This architectural flexibility, which is common in a mainframe environment, is only available on the desktop by using computing coprocessors.

State of the Market
The table lists the manufacturers that are currently shipping a coprocessor product. In order to exclude nonprogrammable systems (such as VGA video cards), table entries are restricted to systems that have a high-speed CPU, will execute user-written programs, and are targeted at desktop platforms.

Some of these systems are intended for specific applications, such as graphics. Boards from Apple, Hewlett-Packard, IBM, and NEC fall in this category. Others are intended as array processors. They don't usually allow direct and easy porting of applications software, as the array processor interacts with the host via a series of subroutines that are linked with main software running on the host. Array processors used to be very popular, but most AP manufacturers now offer alternative, more comprehensive, software-support packages.

continued

COMPUTING COPROCESSORS

The coprocessors listed here fall into three categories. The ones marked "AP" are devoted to array processing. The entries from Apple, IBM, Hewlett-Packard, and NEC are graphics coprocessors. The remainder are general-purpose coprocessing systems.

Company	CPU	Platform	Operating system	Max. RAM (MB per CPU)
Apple	RISC 29000	Mac	QuickDraw	N/A
Avalon	RISC 88000	VAX	VMS	80
CSA	Transputer	PC	DOS	8
Definicon	CISC 68030	PC	DOS	8
	RISC SPARC	PC	DOS	8
	Transputer	PC	DOS	4
HP	Graphics engine	PC	AP	N/A
IDT	RISC MIPS R3000	Mac	Special	64K bytes
IMS	Custom	Mac	Smalltalk	64
	Custom	PC	Smalltalk	64
IBM	RISC i860	PC	AP	N/A
	RISC RT	PC	Unix	N/A
Levco	Transputer	Mac	Mac OS	N/A
Meiko	Transputer	Sun	SunOS	N/A
Mercury	RISC Weitek	PC	DOS	N/A
	RISC Weitek	Sun	SunOS	N/A
Microway	Transputer	PC	DOS	16
NEC	Graphics engine	PC	AP	N/A
Opus	RISC 88000	PC	Unix	16
	CISC 32532	PC	Unix	N/A
Tektronix	RISC 88000	Mac	AP	8
Parsytec	Transputer	PC	DOS	N/A
	Transputer	Mac	Mac OS	N/A
	Transputer	PS/2	DOS	N/A
YARC	CISC 68030	PC	DOS	8
	CISC 68020	PS/2	DOS	4
	Transputer	PC	DOS	16
	RISC 29000	PC	DOS or AP	16
	RISC 29000	Mac	Mac OS or AP	8

The IMS product is also a little unusual. It lets you define the instruction set of the custom processor chip. IMS calls its technology WISC, for writable-instruction-set computer. In addition, IMS offers conventional language-development environments, such as Smalltalk and C++.

The Movement to RISC
One of the startling points visible from even a cursory glance at the table is the number of RISC CPUs that are used in coprocessing systems. The main reason for this is the current higher performance of RISC engines compared to that of conventional complex-instruction-set-computer (CISC) processors. There are several reasons why RISC CPUs are well suited to a coprocessor environment.

First, a coprocessor is a small computing system featuring a tightly coupled RAM subsystem. This happens to be the ideal topology for a RISC computing system. At the high frequencies that RISC processors use, it is very difficult to propagate signals around a circuit board. The bigger the board, the more features it has and the lower its performance. Conversely, the small, tightly coupled coprocessor can often extract peak performance out of a RISC CPU.

Second, RISC carries no requirement for object-level compatibility. Almost all compatibility between coprocessor platforms is achieved at the source code level. Thus, the lack of an applications software base for RISC engines is no particular disadvantage.

In addition, a RISC engine is often less expensive than a CISC CPU of comparable performance. This is especially true for the Advanced Micro Devices (AMD) 29000 and SPARC-based products. Since the cost/performance ratio is often the major selection criterion for a coprocessor engine, the RISC machines usually excel.

RISC vs. CISC
RISC technology has other advantages over CISC that are not specific to coprocessing systems. These days, processors based on RISC define the standards of high-end computing. Although RISC technology itself is not necessarily superior to CISC, a number of factors have combined to position RISC systems at higher performance levels than their CISC counterparts.

First of all, architects of RISC microprocessors began with a clean slate. There were no previous generations of software with which they had to maintain compatibility, and no preconceived notions of how they had to achieve hardware interfacing. This factor is very important.
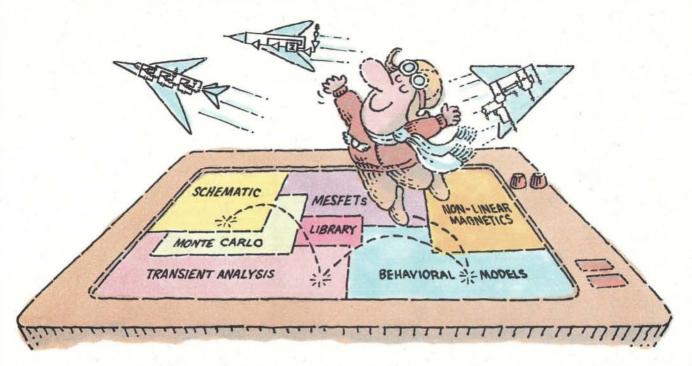
Designers of computers based on a particular family of processors—the Motorola family, for example—are prepared to accept only minor changes in system architecture from one generation to the next. If, for instance, the designers of the 68040 had implemented an external Harvard architecture, with separate external instruction and data memories, they would be in a considerable amount of trouble right now.

Both hardware designers and software engineers would have trouble dealing with such a radical change in 680x0 architecture. Hardware designers would have to redesign system boards, and software developers would have to relink all current 680x0 binaries, for instance, to separate the source code and data address spaces. This would complicate the smooth upward migration path provided by the 680x0 family.

Thus, Motorola implemented an internal Harvard architecture with separate data and instruction buses on the 68040. The data and instruction paths pass through separate memory management units and caches within the chip, but the signals are combined before they are brought out of the chip.

On the other hand, the AMD 29000 RISC designers had no qualms about implementing an external Harvard architecture. Separated external buses seemed clearly superior (and, ultimately, cheaper), so they were used. The designers were not faced with any overriding compatibility issues.

Second, it's easier to design a fast memory system for a modern RISC chip, because the software can maintain closer control of the internal pipeline status

than it can with the CISC chip. This allows the use of "load scheduling" and other optimization techniques to hide the wait states that are introduced when interfacing to real-world memory systems.

Which RISC?

In terms of overall system speed (as defined by the actual speed of operation when running application code), the leading RISC families are currently the MIPS R3000, the AMD 29000, and the Motorola 88000. The R3000 requires too many peripheral and cache chips to be really suitable as a coprocessor engine. Indeed, the only offering is from IDT. This card should probably be used solely as a development platform, as it has only cache RAM and has to go across the Nu-Bus for access to main memory, an unacceptably slow process.

The Motorola 88000 is a very fast CPU, with arguably the fastest single-precision floating-point engine available. It produces its best results on small programs, such as benchmarks, that easily fit within its small 16K-byte data cache. It doesn't perform as well with larger problems primarily because of a minor design flaw in the 88200 cache controller. When there is a cache miss, the whole 4-word cache line must be refilled before execution can continue.

Thus, for programs whose data has good locality of reference, the 88100/88200 chip set turns in a superlative performance (its Whetstone figures are essentially equivalent to those of the Cray-1S). But when it is manipulating sparse matrices, at least 2 words of each cache line are usually never used. The overhead associated with waiting for them to be filled on each cache miss is crippling.

For example, when the 88000 runs a Gauss-Jordan elimination of a 200-by-200 matrix (from the Argonne Labs LINPACK benchmark suite), its Cray-relative performance drops to only 10 percent of what it achieves on Whetstone tests. In fact, if you use enough 88200 cache-controller chips, you can perform the whole LINPACK reduction in the caches and double the performance of the 88000 system.

The AMD 29000 offers a compromise between performance and hardware cost. It uses two totally separated buses for the instruction and data streams, which allows standard DRAM to achieve almost the same performance as static RAM (see "Real-World RISCs," May 1988 BYTE). Although a 29000 system won't outperform an R3000 or an 88000 system on small benchmarks, it will frequently emerge the winner on larger tests or applications code.

Coprocessing Comes of Age

Coprocessor systems have truly come of age. In 1984, when Dave Rand, George Scolaro, and I were defining what we thought could be done with coprocessors and how to do it, we could only dream that one day manufacturers such as IBM, Apple, and Tektronix would offer products using this technology.

Today's coprocessor systems give PCs the computing power of a mainframe. Over the next two years, you will see coprocessors achieve a level of performance currently reserved for supercomputers. Why dream about a Cray? Coprocessor technology will soon put that kind of power into your desktop computer. ∎

Trevor Marshall is a BYTE consulting editor and is president of YARC Systems Corp. (Agoura, CA). You can contact him on BIX as "tmarshall."

Super Sources

Desktop supercomputers rely on a mix of technologies to deliver Cray-like performance in personal computer–size packages. For more information on the developing technologies discussed in this section, contact these companies:

Arche 486 EISA
Arche Technologies, Inc.
48835 Kato Rd.
Fremont, CA 94539
(800) 422-4674
Inquiry 1053.

C-Linda
Hypercomputer Linda
Scientific Computing
Associates, Inc.
246 Church St.
New Haven, CT 06510
(203) 777-7442
Inquiry 1054.

ComputeServer
Torque Computers, Inc.
100 Varick St.
New York, NY 10013
(212) 925-1715
Inquiry 1055.

Connection Machine
Thinking Machines Corp.
245 First St.
Cambridge, MA 02142
(617) 876-1111
Inquiry 1056.

Cray Y-MP
Cray Research, Inc.
608 Second Ave. S
Minneapolis, MN 55402
(612) 333-5889
Inquiry 1057.

i486
i860
EISA chip set
Intel Corp.
3065 Bowers Ave.
Santa Clara, CA 95051
(408) 765-8080
Inquiry 1058.

Intelligent Drive Array
Systempro
Compaq Computer Corp.
20555 FM 149
P.O. Box 692000
Houston, TX 77269
(713) 370-0670
Inquiry 1059.

isPC/860
Intel Scientific Computers
15201 Northwest Greenbrier Pkwy.
Beaverton, OR 97006
(503) 629-7629
Inquiry 1060.

MC486
Apricot Computer plc
Apricot House
111 Hagley Rd.
Edgbaston
Birmingham B16 8LB
UK
44-012-456-1234
Inquiry 1061.

Micro Channel
IBM Corp.
Old Orchard Rd.
Armonk, NY 10504
(914) 765-1900
Inquiry 1062.

MPX
Unix System V/386
Xenix
The Santa Cruz Operation
P.O. Box 1900
Santa Cruz, CA 95061
(408) 425-7222
Inquiry 1063.

Parallel Programming
Support Environment (PPSE)
Oregon Advanced Computing
Institute
19500 Northwest Gibbs Dr.,
Suite 110
Beaverton, OR 97006
(503) 690-1220
Inquiry 1064.

68040
88000
Motorola, Inc.
6501 William Cannon Dr. W
Austin, TX 78735
(512) 891-2000
Inquiry 1052.

SX-X Series
HNSX Supercomputers, Inc.
77 South Bedford St.
Burlington, MA 01803
(617) 270-7600
Inquiry 1065.

386/smp
486/smp
Corollary, Inc.
17881 Cartwright Rd.
Irvine, CA 92714
(714) 250-4040
Inquiry 1051.

Vectra 486 PC
Hewlett-Packard
19091 Pruneridge Ave.
Cupertino, CA 95014
(408) 725-8900
Inquiry 1066.

Z-1000
Z-386/33E
Zenith Data Systems
1000 Milwaukee Ave.
Glenview, IL 60025
(312) 391-8860
Inquiry 1067.

This resource guide lists information sources for the major concepts and products listed in this State of the Art section. Inclusion in the resource guide should not be taken as a BYTE endorsement or recommendation. Likewise, omission from the guide should not be taken negatively. The information here was believed to be accurate at the time of writing, but BYTE cannot be responsible for omissions, errors, or changes that occur after compilation of the guide.

Something just happened to the world's best-selling family of personal computers.

It just got

Six powerful new systems.
Now available are six new PS/2®'s with
MicroChannel™ for those wanting
increased storage, expansion slots and
386™ power.

We've expanded our PS/2 Model
80 line to include two new 20 MHz
models as well as two new cached
25 MHz machines. And two of these
new PS/2 models come with our
lightning-fast 320MB hardfile.

Our PS/2 line has also been
extended with the introduction of two
new Model 65SXs which offer greater
expandability and file capacity along
with an affordable price tag for
386SX™ power.

**New fully upgradable
LaserPrinter E.** The quality
printing and superb paper han-
dling of the IBM LaserPrinter are
now available in the first afford-
able laser printer that's fully
upgradable—the IBM
LaserPrinter E.

It can print at up to 5 pages
per minute and can easily be
upgraded to 10 pages per minute.
What's more, a Postscript®
option is now available for both
the LaserPrinter and the
LaserPrinter E to make your
desktop publishing look its best.

**Portability, power and a
low price.** The award-winning
P70 with *MicroChannel* has a
new 386 16 MHz model. This
new version of our full-function
portable has a 30MB hardfile
and allows you to take your
applications with you wherever
you go.

Innovative software solution
for networking. Connectivity with
your PS/2 is better than ever with t
availability of OS/2® EE V1.2 and
OS/2 LAN Server V1.2.

OS/2 EE is now easier to insta
and gives you ready access to
OS/2 relational data bases and
expanded access to host computers
and Ethernet® networks.

OS/2 LAN Server offers
increased price performance, maki
it one of the lowest cost-per-user
software environments in the
industry.

And to make it all easier, IBM
introduces Easel® 1.1 for OS/2 EE–
the revolutionary development tool.

better.

New high-resolution 14″ display. Excellent resolution, contrast and clarity make the new, affordable IBM 8515 14″ display a pleasure to work with.

Inexpensive new compact disc technologies. Inexpensive large-capacity storage is now yours on any *MicroChannel* PS/2. It's the new IBM CD-ROM and it allows you to store the equivalent of over 400 diskettes of reference material and clip art for applications like desktop publishing, multimedia and more.

New 320MB lightning-fast hardfile. The new 320MB hardfile option, with a lightning-fast access speed of 12.5 MS, comes standard on two of the new Model 80s and is available as an option on all Model 60s, 65SXs and 80s.

Increased expandability with the new Small Computer System Interface (SCSI). An industry-standard SCSI adapter comes with the new 80 and 65SX models and is also available as an option for all your current PS/2s with *MicroChannel*.

Every SCSI busmaster adapter allows you to add up to seven peripherals, so adding an array of scanners, optical disks and new technologies like multimedia and IBM's CD-ROM is no problem.
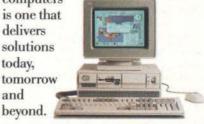
The IBM PS/2 is the world's best-selling family of personal computers. And we haven't stopped there—in fact, we've just made our family even better.

Whether your business is large or small, your investment today in the PS/2 family of hardware and software applications can pay dividends tomorrow. Because the IBM PS/2 family can easily expand with your business and change to meet your needs.

To help you protect your investment, we're offering a wide range of new products. Products that offer more expandability and outstanding price performance. Innovative products that increase the value of your investment in *MicroChannel* and OS/2. Versatile products that offer more solutions than ever, especially for power users and those looking to network.

But more important, we've made these changes in our family so that one thing could stay the same: the IBM PS/2 family of personal computers is one that delivers solutions today, tomorrow and beyond.

How're you going to do it? PS/2 it!

IBM

AROUND THE WORLD IN TEXT DISPLAYS

You think your computer is complicated?
Try typing with 8000 different characters.

Ben Smith

A ccording to the Bible, all the people in the world once spoke a common language. Then, when men of a great city tried to build a tower to the sky, the Lord came down and confused the language and scattered the people. The city came to be known as Babel, after the Hebrew word for "confused."

"Confused" aptly describes today's world, where the existence of thousands of distinct languages creates problems for commerce, politics, and technology—including computers.

Those of us in the English-speaking computer community tend to forget that computing is done all over the world, where functions that we take for granted are not quite so simple. The characters *A* to *Z* are easy to use in computers when compared with Cyrillic, Hebrew, Greek, Arabic, Chinese, and Japanese. Fortunately, today's faster and cheaper processors, memory, and high-resolution screens are providing affordable solutions to the problems of worldwide computing.

First, the Easy Ones

The task of developing systems and software that use non-English characters and words ranges from fairly easy to nearly impossible. Languages that use the same Latin alphabet as English, such as German and French, are the easiest to develop software for. Only a few special characters are required, and you can use these on many word processors by loading international keyboard and printer drivers, or by pressing special key combinations to insert them into documents.

IBM—very much an international company—provides its AIX version of Unix with a set of routines, called *message services*, to help display messages from a program. Developers of programs for the international market can use these facilities to internationalize their applications by externalizing all their messages and help screens. Different message and help tables are built for each language. Modern character-generator tables include extensions to the standard Latin characters for German, French, Portuguese, Spanish, the Scandinavian tongues (including Icelandic), and other European languages.
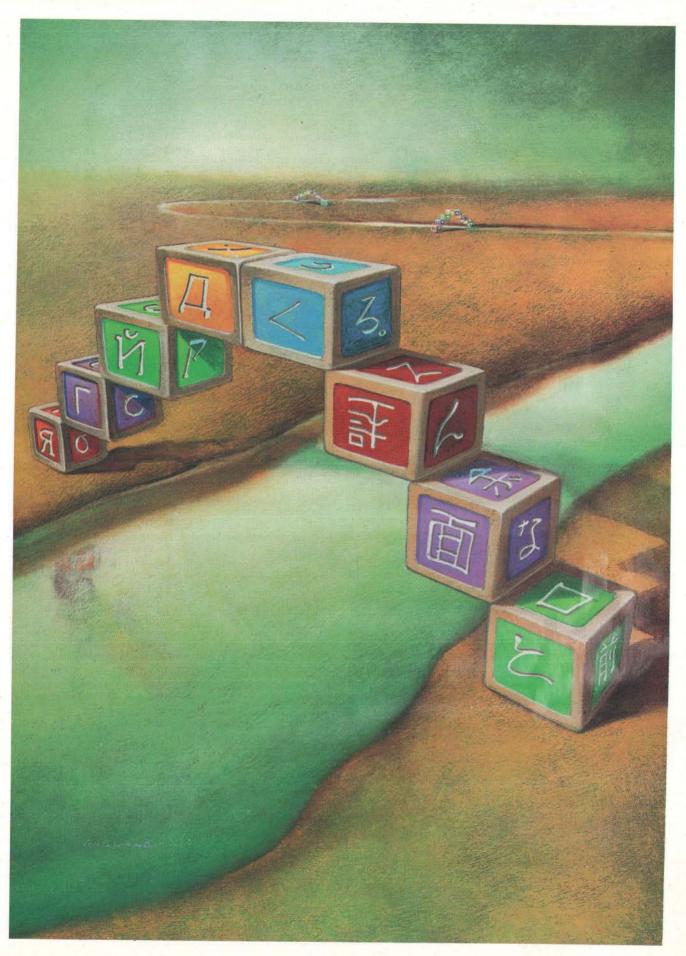
But if the characters are nonroman, say Cyrillic, you need more than just a language-specific message table. You also need to have the characters represented on the keyboard, the screen, and any other output device. Admittedly, the computer knows the characters as only bits and bytes, but the user needs to be able to interact with the computer. There are Cyrillic fonts for screens and printers, as well as Cyrillic key caps. Russian and Slavic writers need to learn only their version of the keyboard and to run versions of software that use their language.

Solutions to Classic Problems

The next major problem of representing foreign languages is how to deal with diacritical marks. This problem stems from the standard way of putting a character on the screen. Each character is stored as an array of bits within a character cell (see figure 1). In IBM PC compatibles and computer terminals, the cell size is dependent on the display hardware. (On bit-mapped systems like the Macintosh, the cell size varies from character to character and from font to font. See the text box "The International Macintosh" on page 266.) When a character is copied from memory to the screen, the entire cell is rewritten, replacing any bits that may have previously been mapped to that area. Thus, there is no way to display diacritics in the same character cell as the base character.

This is an exclusive-OR logic: One or the other character can occupy the space on the screen, but not both at the same time. In order that diacritics be added to the display, there must be a separate character defined for each diacritic and base character (radical) combination. One design that uses this method, IBM's High Function Terminal, has 10-bit characters, which extends the table to 1024 (2^{10}) characters.

A better solution is found by breaking from the standard way of generating characters on the screen, the OR logic, and using AND logic instead; one character *and* another can occupy the same space on the screen at the same time. David Packard Jr. did just this with his design for the Ibycus Scholarly Computer.

You won't find an Ibycus in your local computer store, but you may find one in a college humanities department. It's a special-purpose computer, made to order for classical scholars. It's a word processor and text-searching machine for Latin, Greek, Hebrew, and Coptic writing, including all breathing and diacritical marks and Western European vowels and symbols. Contrasted with the PC's Monochrome Display Adapter 9- by 14-pixel character matrix, the Ibycus uses a matrix 18 pixels wide by 16 pixels high. (The characters are 16 pixels wide; the other 2 pixels separate characters.) The display shows 24 rows of 80 columns.

The Ibycus is usually combined with a CD-ROM device for research and study of the Western classics. The search utilities can quickly find text in any of the supported languages—not a trivial accomplishment.

Closer to personal computer home, Zondervan Electronic Publishing makes a package for biblical and classical scholars. ScriptureFonts lets you enter Greek, Hebrew, and Latin on the same line, display the proper characters on the screen, and print them on your printer. All the characters in the alphabets are mapped to the standard keyboard and displayed using standard character-generation hardware (see photo 1).

Such software must deal with a further complication to internationalization: Not all written texts run left to right and top to bottom. Hebrew text, for example, runs from right to left. Display software must place the cursor on the leftmost column after a carriage return and move the cursor to the right after displaying each character in a stream.

Zondervan has solved many of the cursor-direction and line-wrap problems by using absolute cursor positioning after every key press. (Normally, the system's BIOS automatically moves the cursor to the right after each key press; ScriptureFonts, however, bypasses the BIOS and sends an escape code to the screen after each key press, telling it where to place the cursor.) While this solution works, it adds to the program's overhead and can cause problems; for example, inserting text after the fact confuses the line-wrap macros. For a total solution to all the right-to-left writing differences, the control needs to be at the level of the display device. The Ibycus handles this well because its display is designed to work in any direction.

The Challenges of Arabic

Arabic writing is right-to-left for text and left-to-right for numbers. This complexity is further exacerbated by the number of cases in Arabic. While printed Hebrew has only one case, and Roman and Cyrillic have two, Arabic has four. Arabic also uses stretch character forms, which add emphasis and aesthetics to the written word. So much of the Arabic writing that we are familiar with is actually calligraphic. And while standard typewriter and printing Arabic fonts have existed for decades, the list of characters that are acceptable, even required, varies from region to region. This is not just a minor disagreement—this is a religious issue.

As with roman computing, a standard must be well established for Arabic writing to become common for computing. As yet, no convention exists for Arabic-character computing codes. NCR has a half dozen different character maps for its Arabic terminals and printers. IBM has at least 10.

The problem lies with the displays and printing devices on which Arabic is implemented. First, Arabic characters within a word must be physically connected. An EGA or better display handles this reasonably well, and bit-mapped displays and graphics workstations have no trouble with it.

When it comes to printing, the vertical range of Arabic characters and their diacritical marks is beyond the range of single-pass nine-pin dot-matrix printers. However, 24-pin printers have no difficulty with either the resolution or the management of the Arabic fonts. Laser printers are even more flexible because of their capability for very large character bit maps and the ease with which they build compound characters from basic elements.

Putting aside the religious issues of which characters to codify, the industrial issues of standardizing the codes, and the technical issues of printing and displaying them, you still have the problems of four cases and writing in both directions. You can use standard keyboards with Arabic key caps for text entry if you (or the application you're using) have a way to select the correct form of the letter based on the context. To understand how to do that, you need to know some details of Arabic script.

Figure 1: *On a typical IBM PC–compatible system, characters are constrained to a 9- by 16-pixel cell, making it difficult or impossible to display many international characters. When such a system copies a character from memory to the screen, it completely replaces the previous character, which prevents the addition of diacritical marks, such as those needed for Hebrew and Greek. This means that every possible form of a character must be stored in a separate memory location.*

HOW CHARACTERS ARE FORMED

9 pixels

16 pixels

Photo 1: *ScriptureFonts, a package for students of Western classics and the Bible, lets users enter Greek, Hebrew, and Latin on a single line, display the proper characters on the screen, and print them on a printer. All the characters in the alphabets are mapped to the standard keyboard and displayed using standard character-generation hardware.*

Stretching a Point

The Arabic alphabet consists of 28 basic characters and the stretch character (used for aesthetics and emphasis). With the exception of the letters *alif*, *waw*, and *ya*, all characters are consonants. These three exceptions are actually consonants with attached vowel sounds. The vowel sounds are usually added to Arabic by diacritical marks.

When a character stands alone or is at the end of a word, it is written with a bold stroke. When it occurs in the middle of a word, it is usually joined to the next letter by an upward curved stroke. And when it occurs at the beginning of a word, it is often an abbreviated stroke unless it is capitalized. There are six letters in the Arabic alphabet that can be connected only to the preceding letters. These are cast as delimiters in the scheme described below.

The Arabic characters can be divided into four classes: radical, delimiter, space, and numeric. The radical and delimiter characters can each have four forms (cases): capitalized, end-connective, midword, and word-beginning. Once the characters and forms have been defined in this way, it's not difficult to build contextual-analysis tables and rules. Then the only characters forms that need to be mapped to the keyboard are the base forms (usually word-beginning) of the radicals and their accented versions.

When you enter text on a system that can handle Arabic characters, it's a little surprising to watch the form of surrounding characters change as the computer resolves which form to display—not to mention watching the cursor move to the left (see figure 2). Then, when you enter numerals, the cursor stays in one place and the *numerals* stream off to the left, since numerals run left-to-right.

The Great Frontier: Chinese

Chinese is traditionally written in columns, top-to-bottom, with columns running right-to-left. This problem can be overcome, however, because modern Chinese people have learned to read in Western style, by rows left-to-right. The greatest challenge of computerized Chinese, then, comes from the fact that Chinese characters are not phonetic but rather pictographic and ideographic—that is, they represent images and concepts rather than sounds. The entire Chinese language includes more than 40,000 characters, 8000 of which are commonly used. Try putting all those on your standard keyboard.

Many Chinese characters were developed by combining other characters. For example, the character for *dream* is derived from the character forms for *flower, evening, roof,* and *eyes-while-lying-down*. However, as the base forms are combined, they are often abbreviated to simplify writing or modified to improve the character's readability, writability, or aesthetic value (see figure 3).

Modern Chinese writing was developed with brush and ink. The strokes from a brush naturally have different thicknesses at various places. Uniform lines and curves make reading the characters more difficult, and, in some strokes, impossible. Even though the height of the characters can be kept relatively uniform while using the Western writing direction, character width will vary. Reading Chinese characters with fixed width and height is like trying to understand someone who's speaking in an absolute monotone. Again, the Chinese have learned to cope, this time with the limitations of computer displays.

Chinese computing involves many of the same issues as computing in Arabic: the politics of which characters to include in the basic set, the necessity for high-resolution screens and bit-mapped character displays, printer limitations, and so on. Even if you solve those, you're left with the problem of how to enter the thousands of Chinese characters from a keyboard—a challenge much more complex than that of Arabic.

The first attempts to solve this problem have involved huge, specially designed keyboards with hundreds of keys. Some early designs came from Multitech, a multinational high-tech company headquartered in West Germany. In 1987, Multitech changed its name to Acer and moved its headquarters to Taipei, Taiwan. Suddenly, finding ways to make Chinese computing easier became a necessity.

Besides being a major force in the development of Chinese MS-DOS, Acer has also implemented methods for entering Chinese characters with a standard AT-style keyboard—though "selecting" may be a better word than "entering." Acer provides two ways of selecting a character from the expandable character database in its Chinese system. The first selection method, involving phonetics, migrated from Japan along with many Japanese phonetic characters on the keyboard.

Using 42 key combinations, the user enters the consonants,
continued

ARABIC: A MOVING DISPLAY

Display result	Key press

Figure 2: *Using NCR's Decision Mate V system, a user might type an Arabic word by pressing the characters shown above. After each key press, the cursor moves to the left. When the user presses a space as the last character, the system adds a bold stroke to denote the end of a word.*

CHINESE: A THOUSAND PICTURES

Flower

Eyes-while-lying-down

Roof

Evening
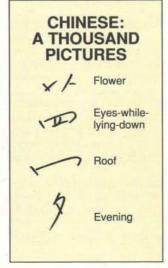
Figure 3: *Chinese characters are based on images and concepts instead of phonetics. Often, these images are combined to form new characters. The Chinese character for* dream *comes from the character forms for* flower, evening, roof, *and* eyes-while-lying-down.

The International Macintosh

Many systems specifications for internationalization are currently little more than just words on paper. But the Macintosh has everything in place and working. The key lies in the Script Manager, a set of general text-manipulation routines for Latin-based alphabets and nonroman writing systems such as Japanese and Arabic. By merely redefining a few resources, an application can take on the appearance of a different language.

Resources underlie Mac applications and system operations. They can be added, subtracted, or modified without any concern by the application, yet their encapsulated information determines the attributes of operations and events (e.g., fonts, menus, and cursor type and motion). Resources in the system file are globally available; resources in the application file are only available while that application is running.

To create an international version of a program, an application developer needs to define messages in the target language. He or she may also have to provide the fonts, although fonts are available for nearly every writing method. Due to the bit-mapped nature of the Macintosh display, the size of characters may vary from character to character and from font to font—a necessity for displaying many international characters (see photo A).

Use of the Script Manager and associated resources greatly facilitates the development of multilingual applications. The Macintosh resources take care of all the details of selecting the correct
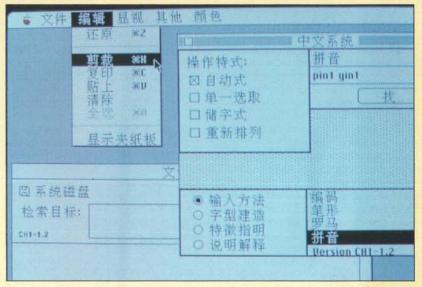
Photo A: *Due to the bit-mapped nature of the Macintosh display, the size of characters may vary from character to character and from font to font—a necessity for displaying many international characters.*

font, displaying it, and controlling the direction of the writing—left-to-right, right-to-left, even up-and-down. Meanwhile, the keyboard is mapped appropriately. The Script Manager defines 32 different scripts, including Japanese, Chinese, Korean, six other Eastern scripts, 10 different East Indian scripts, and even Ethiopian. Different scripts can be combined on the same line, usually a simple problem when the direction of writing differs between scripts.

The Macintosh International Utilities Package works in conjunction with the

Script Manager to complete the task of internationalization: date, time, monetary, and decimal format, as well as sorting where characters may be equal at one level of sorting but unequal at another.

The tragic flaw is that the Macintosh is proprietary and limited in scope. The operating system doesn't fill the needs of non-Mac users. Nonetheless, despite Apple's attempt to limit the use of its ideas, its solutions to human-interface problems have influenced MS-DOS, OS/2, and Unix.

---

vowels, and pitches to describe the sound of the character. Since many different characters can sound the same, the user might have to select the correct character from a list of homonyms. For example, Chinese has three "ying" characters; they look considerably different, but one is for an eagle, another is for a parrot, and the third is for a greenfinch.

Although the process of entering Chinese characters phonetically sounds slow and complex, it actually isn't far removed from entering text in roman-based computing, since the Latin alphabet is phonetic. The difficulty for Chinese computer users is in learning a totally foreign way of thinking about their writing.

The other method of Chinese character selection uses a combination of one to five character elements (out of 25 different elements) to build the character. This is similar in concept to actually drawing the character. Acer's name for the method is Changjei, literally translated "characters of Chang, originator of writing."

While this method of building characters may seem simple,

it took a tremendous amount of work to distill Chinese writing to this structure. Mathematically, more than 10 million characters can be described in this way, although not all possible combinations of the 25 Changjei elements generate valid Chinese characters—just as all combinations of five Roman characters do not necessarily form valid English words.

No Western font or character table is designed to hold even a few thousand characters, let alone the many thousands of Chinese characters. Acer uses TCA Code (Taipei Computer Association Recommended Internal Code) and extended memory to address a 2-byte (16-bit) character-generation code. Once the system has determined which character to use (using one of the two methods described above), it copies the character-generation operations from a table entry to the character generator. There are actually three tables: the first holds the data for the 5402 most commonly used characters; the second holds data for 7650 less frequently used characters; and the third table holds data for 5640 user-defined characters.

Figure 4: *Japanese text, while using a much smaller set of characters than Chinese, can contain three different kinds of writing: the formal kanji, the phonetic katakana, and the informal hiragana—not to mention roman characters for words and phrases that are untranslatable. The English translation is given below the Japanese text.*

JAPANESE: THREE IN ONE

かなり以前のことだが，CP/M 上のデータベース・プログラムである Vulcan というソフトを批評したことがある。そのとき "頭にくるが傑作" という評価を下した。その意味は，強く希望していたことを実現してくれるプログラムであるが，その操作のやり方が

嫌になるほど面倒ということだ。アルゴリズムはとっても優れているが，それに比べユーザ・インタフェースがかなり見劣りしている。余談ではあるが，これはその後dBASE II となった。George Tate は私の批評を読んでVulcan を買収した。

Long ago, I reviewed a CP/M database program called Vulcan that caused me to invent the rating "infuriatingly excellent"; meaning that the program did something I very much wanted done, but I had to fight it every step of the way.

The algorithms were far better than the user interface. Incidentally, that was the first review of the program that became dBASE II. George Tate once told me he bought Vulcan (and renamed it) as a result of my review.

☐ Kanji   ☐ Katakana   ☐ Hiragana

Japanese High Technology

Japanese writing uses a much smaller set of characters than Chinese, but it is more complex because it includes three different *kinds* of writing (see figure 4). The first writing style is kanji (writing of the Kan dynasty of China). This is the traditional and formal style of writing. The second kind of writing is katakana, a phonetic alphabet used to spell out words that are foreign to Japanese. The third kind of writing is hiragana, an informal, abbreviated style. Japanese computers have to display a mixture of kanji, katakana, hiragana, and ASCII roman.

Japanese characters are usually entered phonetically, often spelled out using the anglicized version of the character. For example, to enter the character for *sushi*, the user enters S-U-S-H-I. Simple for us, but not intuitive for Japanese. Another degree of sophistication hides behind all of this: The pronunciation of a word or character depends on where and how it is used in a statement. As with Arabic, characters change on either side of the cursor as the user enters text. If the word isn't resolved to kanji or hiragana, it will end up either katakana or roman.

Because of the computational power and sophisticated programming required to select and generate Japanese characters, computerization of the average Japanese office didn't take off until 1980. Prior to that, all office work was done by hand. There has never been a successful Japanese typewriter.

The Japanese interest in AI grew out of the work on solving the data-entry problems of their writing, likely the most complex of all. Now Japan has become a great resource for innovative programming as well as computer hardware. Although the Japanese started computerizing much later than the Western world, the lessons they learned on the way have made them leaders in sophisticated systems.

A Unified World Theory

Each writing method carries some special problems for application programs. For example, how do you sort textual data when an alphabet has more than one sorting order? How do you edit Chinese writing? Do you remove the entire character and

start again, or can you edit the character? Sorting Chinese names is simplified by the fact that there are only about 200 Chinese surnames. The Japanese use a great deal of Chinese writing, and there are far more surnames in Japan. So Japanese names are sorted phonetically.

If you think Japanese computing has problems, consider India, which has nearly 20 different methods of writing, besides classical Sanskrit, Devanagari, and Arabic. Countless other writing methods exist in Asia and Africa. Africa alone has nearly a thousand languages—many with no written form.

Still, there are now affordable solutions to the technical problems of computing for the most widely used forms of non-roman writing: Eastern European (Cyrillic), Hebrew, Arabic, Chinese, and Japanese. But great challenges still remain. Perhaps the greatest challenge involves the issue of standards. Some promising signs exist: Many modern standards require programs to have an international version before they can be certified. ANSII Standard C defines a multibyte character type and specifies that all string functions operate on the structure in a sensible manner. The IEEE Posix standards for Unix include specifications for internationalization.

These are steps in the right direction. Structures that allow for more information per character position must exist below the application program, at the level of the compiler and operating system. Unfortunately, very few operating system and language implementers have done more that acknowledge the problem and agree to work on it "very soon now."

Until the problem is attacked at that level, international users will be forced to make do with ad hoc solutions. Not that we need more standards; the existence of YAT (Yet Another Standard) would only make the problem worse. We just need to change our mind-set from "I can do it better" to "We can do it together." ∎

Ben Smith, a technical editor for BYTE, formerly helped to develop a computerized form of Arabic. He can be reached on BIX as "bensmith."

Elegance*

THE NORTHGATE 386 33 MHz ELEGANCE SYSTEM

*"**1a.** Refinement and grace in movement, appearance or manners.

b. Tasteful opulence in form, decoration or presentation.

2. Something that is elegant." (American Heritage Dictionary)

You said it.

Since at least three of the medals garnered by General Northgate were for our highly
acclaimed Elegance series (triple Editors' Choice awards from PC Magazine, for example),
we thought we'd show you what the machine looks like and give you a few specs
in case you might be inclined to buy this elegantly designed state-of-the-art computer.

First of all, to photograph an elegant machine...you need an elegant
photographer. So we went to the most highly respected lensman in Hollywood:
Harry Langdon. He normally lights and shoots such famous faces as Linda Evans,
Victoria Principal, Cher, Arnold Schwarzenegger and Diana Ross, to name just a few.
So shooting a different pretty face like a Northgate Computer is all in a day's work for
Harry. For one thing, Northgate doesn't need a hairdresser.

The Elegance series to your left comes in three versions: The 20 and 25
MHz models, and our top of the line/highest performance Elegance: the 33 MHz which
earned a 9.1 rating in Infoworld.

The price of the complete 386™ 33 MHz Elegance system pictured: $5995.00.
(Of course, you can buy a fully configured Elegance system starting at $3395.00.)

The only thing more impressive than the price is the tech support which
comes along with our usual 1 year parts and labor warranty. Did we mention that we
also ship you replacement parts, should they be needed, overnight at our expense?

And that our tech support people are actually here for you 24 hours a day,
every day, every night?

Put us to the test: give us a call some night at midnight. Or 3:00 AM.
Now to the specs:

Technobabble*

*Seemingly endless Technical specifications, without which you aren't about to buy
this machine no matter how many awards it's won.*

*Not only that...would this look like a real computer ad without them? Probably not.
Here they are:*

STANDARD CONFIGURATION
- 33 MHz Intel 80386 processor.
- 150 Mb ESDI Hard Drive 16MS.
- 4 Mb Ram.
- 64 K SRam Read-Writeback Cache
  (Optional 256 K Cache available)
- Zero wait state performance.
- 1.2 Mb—5.25" floppy drive (also reads,
  writes and formats low density diskettes).
- 1.44 Mb—3.5" floppy drive (also reads,
  writes and formats low density diskettes).
- Eight expansion card slots
- 80387 or Weitek coprocessor support

- Two serial ports and one parallel port (two
  with monographics systems)
- 14" Multifrequency VGA Color Monitor
  1024x768 resolution with 16 bit controller.
- Sleek new Elegance 7 drive-bay custom
  vertical cabinet pictured. (Desktop style
  optional, save $150.00.)
- Exclusive award winning Omnikey
  Keyboard.
24 Hr 7-days a week tech support: unlimited
toll free phone.
Additional Technobabble available upon
request

Now Northgate™ Gives You Credi

Pay for your Northgate system with budget-easing monthly installments.

Owning any Northgate system is easy and affordable when you use Northgate's Big 'N' credit card. Own the computer system voted Best Buy by *Computer Shopper* without straining your budget or tying up your major credit cards. When you Charge-it on Big 'N' your payments fit your budget every month.

Simply fill out the Big 'N' information form at the right and send it to Northgate Computer Systems, P.O. Box 41000, Minneapolis, MN 55441. We'll do the rest. Better yet, call our toll free number and talk to a Big 'N' representative.

After your credit is approved, one of our expert sales representatives will help you design the Northgate system that's just right for you. Best of all, your system will be shipped promptly.

Northgate also offers flexible long term leasing plans for your business. You can choose the plan that best suits your needs, with up to 5 year terms available.

Charge-it on Northgate's Big 'N' credit card! It's EASY to qualify, CONVENIENT to use with 24 hours a day, 7 days a week service, and assures FAST delivery after you order.

Northgate, the Northgate Big 'N' and SlimLine are registered trademarks of Northgate Computer Systems, Inc. ©Northgate Computer Systems, Inc. 1990. All rights reserved.

BIG 'N' REQUEST FOR CREDIT

NOTICE TO WISCONSIN RESIDENTS
You must disclose your marital status.
☐ married
☐ unmarried
☐ legally separated

A married person may apply for individual credit. I am applying for (please check appropriate box):
☐ **JOINT CREDIT** with another person. Complete entire application.
☐ **INDIVIDUAL CREDIT** but rely on income or assets of another person as a basis for repaying the credit requested. Complete entire application.
☐ **INDIVIDUAL CREDIT.** Complete sections ''a'' and ''b'' only.

Please complete all appropriate sections, providing at least two year's residence and employment history. This will enable your application to be processed as quickly as possible. If you are self-employed, please be sure to complete section ''D'' on back.

Applicants must be 18 years of age or older.

a. Your Personal Information
Requested Line of Credit $

Your Name: First / Initial / Last / Date of Birth: Mo. Day Yr. / Social Security Number:

Present Address Street / Apt. # / City / State / Zip / Home Phone ( )

Date of Residence: Month / Year / Monthly Payment: $ / Buy ☐ Rent ☐ Other ☐

Previous Address: / Dates of Residence: From To

Your Employer: (If self-employed, see rear panel) / Date of Employment: Mo. Yr. / Position: / Monthly Income: Gross $ Net $

Employer's Address: Street / City / State / Business Phone: ( )

Previous Employer: / Address: / Dates of Employment: From To

Income from alimony, child support or separate maintenance payments need not be disclosed if you do not wish to have it considered as a basis for repaying this obligation. / Other Income: / I have received since: (Date) / Monthly Income: Gross $ Net $

Name and Address of Nearest Relative Not Living With You: / Relationship:

b. Credit Information
Include joint applicant's information, if joint account requested.

	Bank Name	Address	
Bank Account:			☐ Checking ☐ Savings
Bank Account:			☐ Checking ☐ Savings
			Payment / Balance
Bank Loan Reference:			
Bank Card Reference: ☐ VISA ☐ MasterCard			
Other Credit Card Reference			
Other Credit References:	Account No:	Expires:	

Driver's License No. / State: / Expires:

c. Joint Applicant's Personal Information
*If you are a married Wisconsin applicant, you must provide your spouse's information below, even though your spouse may not be signing this contract.

Joint Applicant's Name: / First / Initial / Last / Date of Birth: Mo. Day Yr. / Social Security Number:

Address Street / Apt. # / City – State / Zip / Date of Residence: Mo. Yr. / Home Phone ( )

Employer: / Date of Employment: Mo. Yr. / Position: / Monthly Income: Gross $ Net $

Employer's Address: Street / City / State / Business Phone: ( )

d. Self-Employed Information
Complete this section only if you are self-employed.

Business Name: / ☐ Proprietorship ☐ Corporation ☐ Partnership

Business Address: / Business Telephone: ( )

Description of Business: / Your Position: / In Business Since:

Your annual income from business: / Business' annual income: / (gross) (net)

You must provide at least one of the following:
1. Business Bank / Name / Telephone ( ) / Personal Banker's Name:
2. Accountant's Name / Telephone ( )
3. Financial statement on business attached.

BYI590

PHONE NORTHGATE NOW
24 HOURS A DAY, 7 DAYS A WEEK.
800-548-1993

NORTHGATE COMPUTER SYSTEMS
P.O. Box 41000, Minneapolis, MN 55441

Circle 179 on Reader Service Card

POWER TO THE PORTABLES

New batteries will store more power and allow laptops to run longer

Andy Reinhardt

I n 1981, Adam Osborne introduced not only a new computer, but a new *kind* of computer: the portable. The original Osborne I, like its early offspring, was the size of a suitcase and was meant to be ported from home to office. It was definitely not a laptop machine—unless you had a lap like the Incredible Hulk—and its AC power cord meant that you had to tether it to a wall outlet when in use.

When Tandy introduced the AC/battery-operated Model 100 in 1983, the portable family tree split into two branches: the luggables and the laptops. The trade-offs between these two types of portable computers involved power, in one form or another. If you needed a small machine that you could take on the road and run off batteries, then you bought a laptop and settled for limited features. If you needed a high-powered CPU, memory, disk drives, screen, and software, you were stuck with a suitcase-size portable.
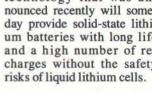
Since those days, the two types of portables have converged, thanks, in part, to smaller, low-power components. Still, a major limiting factor in laptop design has been the humble battery, which has often dictated what you could do with your laptop and how long you could do it.

Experts have long predicted that better batteries for laptop computers were just around the corner. The time is right. Despite big improvements in power management, neither customers nor manufacturers are entirely happy with the current options, especially finicky nickel-cadmium (NiCd) batteries. Actually, every type of portable battery now in use has drawbacks.
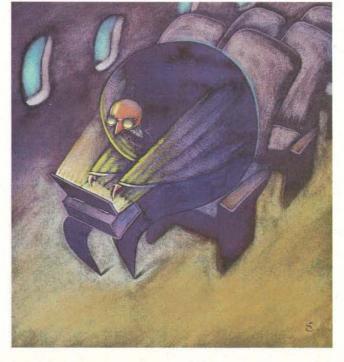
Poised to arrive are new batteries that promise to solve many of the current crop's problems. Nickel-hydride (NiH$_2$) batteries, for instance, offer higher power capacity and thus longer operating life. Air-breathing zinc alloy batteries, similar to those used in hearing aids, offer longer life, lower weight, and fewer hazards. And a new technology that was announced recently will someday provide solid-state lithium batteries with long life and a high number of recharges without the safety risks of liquid lithium cells.

The Trouble with Nickel Cadmium

Rechargeable NiCd batteries are widely used in today's laptops, including models from Compaq, Toshiba, Zenith, and NEC, yet their reputation is not very favorable. Aside from environmental concerns regarding their disposal (cadmium is highly toxic), NiCd batteries are a hassle. Even with sophisticated power management, they typically last only 3 to 4 hours before they need to be recharged, and they can only be

continued

Figure 1: *At one end, the Aerobic Power battery has a fairly typical zinc alloy anode. The cathode, however, features an air-breathing carbon membrane that extracts oxygen from the atmosphere. Oxygen and water are broken down into negatively charged hydroxide ions (OH⁻), which bond with positively charged zinc atoms, releasing water and electricity.*

recharged a few hundred times before they must be replaced.

The greatest controversy about NiCd batteries concerns the recharging process. The energy output of NiCds is relatively stable, dropping off only as they near depletion. But if improperly recharged, NiCds suffer a well-known memory effect known as *voltage depression*. This means that the battery "remembers" how much power was left when you recharged it. When you use your machine and the battery's power falls below that level, the voltage drops 10 percent. This situation can fool the computer's power management circuitry into thinking the battery is spent.

To avoid this problem, vendors counsel users to fully discharge the batteries before recharging them and not to charge batteries in mid-cycle. However, many designers now discount the memory effect, saying that good power management design eliminates it completely. Richard Stickel of Gates Energy Products blames most of the problems associated with NiCd voltage depression on poor power engineering. For example, he says, rather than using 10 1.2-volt NiCds for a 12-V system, designers should use 11.

In the last year, many manufacturers have reworked older technologies in innovative ways to avoid the problems of NiCds. Dynabook Technologies used a dry lead-acid battery for its portable, while Apple chose a lead-acid gel battery for the Macintosh Portable. A growing number of pocket-size personal computers, including the Poqet, Psion MC series, and Atari Portfolio, use off-the-shelf AA alkalines.

But none of these alternatives to NiCd is an ideal solution. Lead acids are as heavy and expensive as NiCds, and their power decays at a more linear rate, putting big demands on the power management system. AA cells are relatively cheap, but they can't be recharged. And, under some circumstances, all these batteries can explode or leak.

Nickel Hydride

Last summer, when personal computer pioneer George Morrow announced his intention to market a laptop sometime in the future, he said that the machine would use NiH_2 batteries, an emerging alternative to NiCds. Morrow said that his nickel hydrides are twice as efficient as traditional NiCds.

Now, battery maker Gates Energy Products is commercializing rechargeable nickel hydrides. Gates says that it intends to

produce samples this year and volume quantities in 1991. According to the company, NiH$_2$ batteries able to replace the NiCds in existing portables will have 50 percent higher capacity, which could boost operating time by as much as 50 percent.

The only problem with nickel hydrides is that they have a higher internal impedance than NiCds. Laptop makers won't be able to count on them to run computers that both consume more power *and* last longer between charges. To best exploit the higher capacity of nickel hydrides—that is, to enjoy longer operating times—laptop designers will have to make computers that use the same amount of or less power.

Stickel says that NiH$_2$ batteries will suffer less from voltage depression than today's best NiCds and that you should be able to recharge them nearly as many times. Current prototypes, however, do not perform this well. NiH$_2$ batteries will cost somewhat more than NiCds, although exact pricing hasn't been established. They are, however, environmentally preferable to NiCds because they don't contain cadmium.

The "Aerobic" Battery

One new development has come from an unexpected source: A small company called Dreisbach ElectroMotive, Inc., or DEMI, recently demonstrated a zinc battery that breathes airborne oxygen for use in its internal chemical reaction. This rechargeable battery, dubbed Aerobic Power, offers lighter weight and greatly improved run time over NiCds and will cost about the same in dollars per hour of operation, DEMI says.

The Aerobic Power battery has a zinc alloy anode and an air-breathing cathode—a carbon membrane that extracts oxygen from the atmosphere (see figure 1). In operation, the Aerobic Power requires a ducted air supply and some sort of airtight shield to shut it down when not in use.

The battery reaction draws oxygen through the gas-permeable membrane into an alkaline gel electrolyte. Oxygen and water molecules in the electrolyte break down into oxygen and hydrogen atoms that pair to form negatively charged hydroxide ions. These ions, in turn, bond at the anode with positively charged zinc ions to form harmless zinc oxide. Water and electricity are released in the process, and each cell produces about 1.2 V of potential. The excess water is absorbed back into the electrolyte.

One of the major breakthroughs of the Aerobic Power battery is its rechargeability. Other companies have attempted without success to design rechargeable zinc-air batteries, says DEMI's Len Danczyk. DEMI's process is proprietary, but basically it involves reversing the reaction. Water from the electrolyte and electricity supplied by a power source are applied at the anode, producing zinc atoms and free hydroxide ions that escape through the cathode membrane as oxygen.

Among this battery's other breakthroughs are its high power output—up to 1000 times that of a typical hearing aid battery—and its low weight-to-power ratio. Zinc-air batteries are lighter than others because the oxygen is obtained from the air, rather than from metal compounds such as the nickel hydroxide or

ENERGY DENSITY IN WATT-HOURS PER KILOGRAM

Legend: Nickel cadmium, Lead acid, Nickel hydride, Alkaline, Solid-state lithium, Zinc air
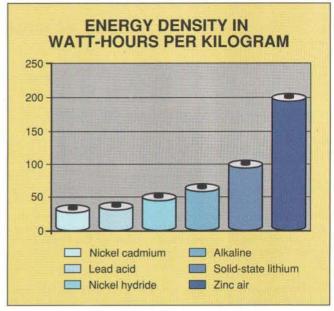
Figure 2: *As this graph of power output shows, rechargeable nickel-cadmium batteries (currently favored for portables) and good old alkaline batteries may face stiff competition from zinc-air, nickel-hydride, and solid-state lithium batteries, which offer combinations of rechargeability, higher power, and longer operating life.*

manganese dioxide typically found in other batteries. A single cell weighing only half as much as a NiCd cell can produce power four to five times longer, DEMI says. This feature gives the Aerobic Power a rating of 200 watt-hours per kilogram, versus roughly 30 Wh/kg for NiCd and lead acid (see figure 2). Thus, a laptop battery could last for an entire 8-hour workday.

A key problem with zinc-air batteries has been that dendrites of zinc oxide (tree-like marks made by one mineral crystallizing in another) grew on the anode and could eventually short-circuit the battery. DEMI says that its process combats this zinc oxide buildup but won't disclose details on the technology.

A drawback of this battery is that the maximum number of recharge cycles is less than that for NiCds. DEMI currently rates the battery for 25 to 30 cycles; it expects 40 to 60 when the technology matures. But because its raw materials cost less, the Aerobic Power battery will sell for roughly $50, as opposed to from $80 to several hundred dollars for NiCds and lead-acid cells. In the end, the cost per hour of operation is equal, DEMI says. Another potential problem, according to Bill Jergens of Texas Instruments, is that once they are exposed to oxygen, zinc-air batteries tend to degrade no matter how carefully they're sealed.

SOLID LITHIUM: FLEXIBLE POWER

Figure 3: *In solid lithium batteries, a thin film of disulfide polymers (acting as a cathode) is layered between films of lithium (the anode). When the lithium oxidizes, it releases electrons that cause the sulfur compound to depolymerize, giving off electrochemical energy. Much more stable than conventional lithium batteries, this lithium "sandwich" could be made into "coins" (for use in cameras and watches) or flat bricks (such as the batteries used in today's laptops), or rolled up to create various sizes of cylindrical batteries.*

Pilot production is expected this year, with volume laptop battery production slated for early 1991. Incorporation of the Aerobic Power battery into commercial laptops will require higher-power recharge circuitry, changes to power management schemes, and packaging modifications, including the addition of air ducts and airtight access doors.

Laboratory Lithium

Although George Morrow was enthusiastic about nickel hydride, he noted that rechargeable lithium batteries would be even better. Unfortunately, such batteries aren't manufactured because of their potential explosiveness. Currently, the only available commercial lithium cells are nonrechargeable "coin" batteries used in some cameras, watches, and calculators.

But there's hope on the horizon. Scientists at the University of California's Lawrence Berkeley Laboratory recently announced that they have developed a new lithium battery that uses advanced thin-film technology and is based entirely on solid materials. This battery should deliver higher power, more recharge cycles, and a longer shelf life (at a lower price and weight) than any of the commercial batteries that are now available. And, say its developers, unlike lithium batteries with a liquid electrolyte, the new batteries won't leak or explode when exposed to heat.

This solid-state, thin-film lithium cell exploits the organic sulfur-sulfur bond that is essential in chaining together amino acids into proteins. To use this bond innovatively, the scientists developed a breakthrough cathode material—a thin film of disulfide polymers that contains long chains of molecules joined by these sulfur-sulfur bonds (see figure 3).

The cathode is layered between thin films of a lithium anode. When the lithium oxidizes, it releases electrons that cause the sulfur compound to depolymerize, giving off electrochemical energy. Electricity reverses the process, causing the sulfur bonds to rejoin into polymers. According to the developers, no battery has ever before used the energy of sulfur-sulfur chemical bonds.

Solid-state lithium batteries should be able to operate in both room-temperature and high-temperature applications (80°C to 100°C). Room-temperature versions for uses such as laptops could be available in three to four years. High-temperature versions may someday have applications in electric automobiles.

This new technology should result in many ways to configure lithium batteries. Manufacturers could start with sandwiches of thin films and cut them into coins similar to the current lithium batteries. They could layer them into flat bricks such as the batteries used in today's laptops. Or they could roll them up to create the various common sizes of cylindrical batteries. The solid-state lithium coins could deliver about seven times the watt-hours per kilogram of today's liquid lithium cells. Room-temperature cylindrical batteries could deliver about four times the watt-hours per kilogram of a NiCd.

According to its developers, this new type of lithium battery can be recharged 100 times with minimal loss of energy capacity. Some prototypes have been successfully recharged as many as 350 times. In addition, developers believe that the raw materials for the new battery will cost less than those required for current batteries, which would result in less expensive replacements. That this type of battery contains no toxic materials should minimize disposal problems. And one benefit of these new devices that should greatly interest laptop toters (with one arm longer than the other) is that because the polymer in lithium batteries is lighter than a metal cathode, the batteries will weigh less than NiCds.

Power Management

Even with the pending arrival of new battery technologies, most observers say that improved battery performance—especially in the immediate future—will depend on the same criteria that are valid today: lower power consumption and better power management.

NiCd technology has matured. Users probably won't see additional capacity improvements of more than 10 percent to 15 percent. To boost laptop operating time, manufacturers will need to use lower-power disk drives and displays, active power management (CPU sleep mode and system suspend/resume), and "passive" advances such as static CMOS ICs and slow refresh DRAM.

Furthermore, some of the high-storage-capacity batteries of the future will achieve their longer operating times in part by doling out power at a lower rate. Since both NiH_2 and zinc-air batteries perform best when the consumption of current is very low, vendors must provide systems that are even less power-hungry than today's models to take full advantage of the new power technologies. ■
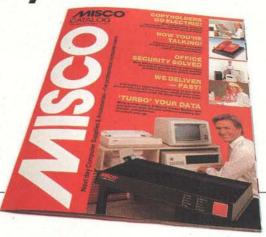
Andy Reinhardt is an associate news editor for BYTE. He can be reached on BIX as "areinhardt."

The Total Scanner Solution

DFI HS-3000 PLUS
HANDY SCANNER

- DESKTOP PUBLISHING
- WORD PROCESSING
- DATA BASE APPLICATIONS
- HIGH LEVEL LANGUAGE INTERFACE
- CAD APPLICATIONS

The Handy Scanner 3000 Plus comes bundled with ZSoft's top-selling PC Paintbrush Plus and DFI's SCAN and ImageTools programs. PC Paintbrush Plus lets you paint, rotate, flip, invert, shrink, tilt, expand and add text into your scanned images. Using SCAN, you can save your images into TIFF, PCX, IMG, MSP, and CUT image files or convert between the formats. Using ImageTools, you can interface the Handy Scanner 3000 Plus with dBase IV, dBase III +, and compatibles such as Cllpper, Foxbase or design custom applications using Microsoft C and QuickBASIC, or Turbo Pascal.

Available as options are DFI's OCR package, which reads mono or proportional spaced typeset characters in English, German, French, and Spanish languages and a fast CAD conversion program that allows you to save and convert scanned images into CAD image formats. At DFI, we offer you the total scanning solution. The Handy Scanner 3000 Plus, a new age in scanning technology.

DFI expands the world of desktop scanning with the new Handy Scanner 3000 Plus. Consolidating all the advanced hardware features you have come to expect from the company that introduced hand scanners to the world, our scanner now gives even sharper images and is even easier to use. Perhaps, most important, DFI gives total scanning support: desktop publishing, word processing, data base, high level language development and CAD applications.

DFI®

SMOOTH VIEWS

Antialiasing lets programs squeeze higher resolution out of VGA displays

Michael A. Covington

emember the first, crude computer graphics? With screen resolutions of about 320 by 200 pixels, and a limited number of colors (in numerous cases, exactly two), early microcomputer graphics forced users and programmers to get creative, both in drawing graphics and in imagining that the image on the screen really looked like what it was supposed to.

That has all changed, thanks to high-resolution screens and technology that allows for millions of colors. Of course, not everyone can afford the latest state-of-the-art video display. But there are ways to get more out of your current display.

One way is a technique known as *antialiasing*, which can give any video display the smoothness—though not the full sharpness—of a display with twice the resolution. Diagonal lines lose their stair-step appearance, and small type becomes more readable. There are trade-offs; a 16-color display is reduced to only two colors, and the displayed image is slightly blurred. But since the human brain is good at compensating for blur, the overall effect is remarkably pleasant and effective. (See photo 1.)

As an example, I'll show how antialiasing doubles the effective resolution of a 640- by 480-pixel VGA display.

How It's Done
The key idea of antialiasing is that a line on the screen cuts a swath across a grid of pixels (see figure 1). It fills some pixels completely and others only partly. To produce an antialiased image, the partly filled pixels are only partly illuminated. For instance, a half-full pixel is displayed at half brightness. This smooths out the jagged appearance of lines and other shapes.

To keep track of how much of a pixel is filled, you divide it into 4 subpixels (see figure 2) and assign 1 bit to each. As it happens, the VGA already assigns 4 bits to each pixel to distinguish 16 colors.

To achieve antialiasing, all you have to do is treat the color bits as subpixels and then translate the colors into appropriate percentages of illumination. The colors corresponding to

binary 0001, 0010, 0100, and 1000 will be 25 percent illuminated; those corresponding to 0011, 0101, 0110, 1001, 1010, and 1100, half illuminated; and so on. The color of each pixel will then show how many of its subpixels are illuminated, but not which ones.

Setting Up the VGA
Procedure SetUpAntialiasing (see listing 1) does all the work. First, it uses SetGraphicsMode to invoke 640 by 480 16-color mode (denoted by hexadecimal 12). SetGraphicsMode also sets the Pascal global variable DirectVideo to false so that all write and writeln statements will use the BIOS; otherwise, the Pascal run-time I/O package would run into problems because it doesn't know the screen is in graphics mode.

Then, SetUpAntialiasing maps VGA colors 0 to 15 onto VGA color registers 0 to 15. At any given time, 64 colors can be defined, but only 16 can be used, so you just use the first 16.

continued

Conventional Antialiased
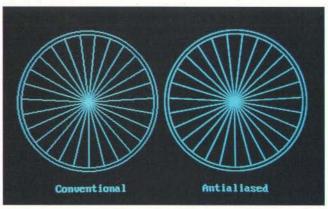
Photo 1: *By smoothing the transitions, antialiasing removes the stair-step appearance of lines and curves.*

ANTIALIASING
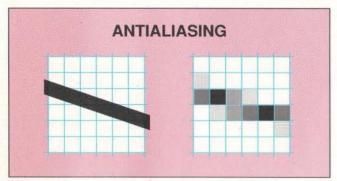
Figure 1: *As a line cuts a swath across a grid of pixels (left), partly filled pixels are partly illuminated (right).*

SUBPIXEL BIT MAPPING
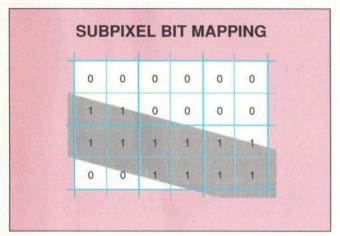
Figure 2: *To keep track of partial filling, each pixel is divided into 4 subpixels. By mapping subpixel bit patterns onto color registers, the display produces the appropriate brightness.*

The next step is to define these 16 colors. Recall that each color number, 0 to 15, is actually a 4-bit binary number, and you are treating the 4 bits as subpixels. So the brightness percentage assigned to each color depends on the number of 1s in its binary representation.

The five constants c0, c1, c2, c3, and c4 give the brightness values for 0 percent, 25 percent, 50 percent, 75 percent, and full illumination. To produce different shades of cyan (black to full bright), you assign the values in the appropriate c variables to both green and blue phosphors, leaving red at 0. You may prefer a yellow-brown (red plus green) or gray-white display instead. You could even get other hues by assigning fractions of the total brightness to red, green, or blue in ways of your own choosing.

Plotting Antialiased Displays

There are two basic operations for plotting a point: PlotSubpixel, which illuminates a single subpixel, and PlotDot, which illuminates 4 adjacent subpixels (see listing 2). A "dot" thus plotted is the size of a full pixel but does not necessarily coincide with one. Plotting with dots results in smooth-edged figures; plotting with individual subpixels does not, because the smoothing effect applies only to the edge of a relatively large area of subpixels.
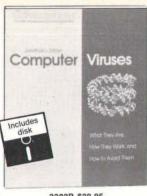
As you might guess, PlotSubpixel sets a single bit in the

Listing 1: *Basic video routines for antialiasing.*

```
program Antialiasing;
uses CRT, DOS;

procedure SetGraphicsMode(mode:integer);
var regs: Registers;
begin
 regs.ax := (mode mod 256);
 intr($10,regs);
 DirectVideo := false;
end;

procedure SetPaletteRegister
                (register,value:integer);
var regs: Registers;
begin
 regs.ax := $1000;
 regs.bx := ((value mod 64) shl 8)
                + (register mod 16);
 intr($10,regs)
end;

procedure SetColorRegister
        (register,red,green,blue:integer);
var regs: Registers;
begin
 regs.ax := $1010;
 regs.bx := register;
 regs.dx := ((red mod 64) shl 8) + lo(regs.dx);
 regs.cx := ((green mod 64) shl 8) + blue;
 intr($10,regs)
end;

procedure SetUpAntialiasing;
const
  c4 = 63;   { Full brightness }
  c3 = 55;   { 3/4  brightness }
  c2 = 48;   { Half brightness }
  c1 = 40;   { 1/4  brightness }
  c0 =  0;   { No   brightness }
var i: integer;
begin
  SetGraphicsMode($12);
  for i:=0 to 15 do SetPaletteRegister(i,i);
  SetColorRegister( 0,0,c0,c0); { for 0000 }
  SetColorRegister( 1,0,c1,c1); { for 0001 }
  ...likewise for all other 4-bit patterns...
  SetColorRegister(15,0,c4,c4); { for 1111 }
end;
```

Listing 2: *Subpixel plotting routines.*

```
procedure PlotSubpixel(x,y:integer);
  { Illuminate a subpixel. }
  { Coordinates are 0..1279, 0..959. }
var
  offset: word;
  plane, planemask, oldval: byte;
begin
  { Select appropriate memory plane }
  plane := (y and 1)+(y and 1)+(x and 1);
  planemask := 1 shl plane;
  port[$3CE] := 4;
  port[$3CF] := plane;    { for reading }
  port[$3C4] := 2;
  port[$3C5] := planemask; { for writing }
  { Set value in memory }
  x := x div 2;
  y := y div 2;
  offset := word(y) * 80 + x div 8;
  oldval := mem[$A000:offset];
  mem[$A000:offset] := oldval or
                    ($80 shr (x mod 8));
end;

procedure PlotDot(x,y:integer);
  { A dot is the size of a full pixel }
  { but need not coincide with one. }
begin
 PlotSubpixel(x,y);
 PlotSubpixel(x,y+1);
 PlotSubpixel(x+1,y);
 PlotSubpixel(x+1,y+1)
end;
```

Bresenham's Line and Circle Algorithms

In 1965, J. E. Bresenham devised the standard way to plot lines on a computer. Bresenham's line-drawing algorithm applies to lines that are within 45 degrees of horizontal. You can bring steeper lines into this range by reversing the roles of x and y.

For the purpose of this explanation, assume that lines slope up to the right (you can deal with negative slopes by reversing a sign). The algorithm starts at one end of the line and moves horizontally to the other end, plotting a pixel in each column. At times the plotted line stays on the same row of pixels, but if the line has any slope that can be represented on the screen, it eventually moves up to the next row in order to follow the slope (see figure A). The problem is how to determine when to move up to the next row.

You have to keep track of the true height of the line. Every time you increment x (move 1 pixel column), y should increase by $(y_2-y_1)/(x_2-x_1)$. Actually, since a pixel row is an incremental element, the row number either stays the same or increases by an integer. If you store the discrepancy in the variable e and add $(y_2-y_1)/(x_2-x_1)$ to it with each new column, then you can advance to the next row when e reaches 0.5. Whenever you move to the next row, you subtract 1 from e.

Bresenham eliminates the slow floating-point operations in the following

LINE ALGORITHM
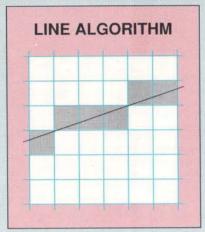
Figure A: *Bresenham's line algorithm displays a line by determining which pixels are appropriate. The math is done with integers, thus avoiding the overhead of floating-point operations.*

CIRCLE ALGORITHM

Figure B: *Following a circle requires vertical, diagonal, and horizontal moves.*

way: If you start e at 0.5 and test for 1, you can use only integers by multiplying the variables and tests by $2(x_2-x_1)$. The sign of e doesn't change, so the comparisons still work the same way. Eliminating the floating-point calculations and tests speeds up the operation.

Bresenham also devised an efficient way to plot circles. Again, the key is to move incrementally and choose pixels one at a time. The center of the circle is

treated as the origin, and plotting starts at $x=0$, $y=r$. The question is then whether to get to the next pixel by moving vertically, horizontally, or diagonally (see figure B).

In making its decision, the algorithm thinks about moving diagonally to $(x+1, y-1)$ and asks how close to the circle it would end up. There are three possibilities. If it would land very close to the circle, it goes ahead and moves di-
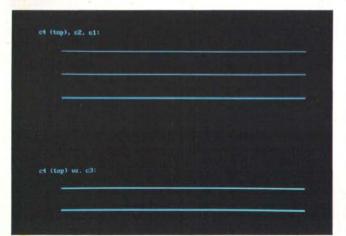
Photo 2: *When you display the test pattern, the width of lines is inversely proportional to the desired brightness; if the lines look alike from a moderate distance, then the balance between width and brightness is correct.*

VGA's video memory. This is a surprisingly complex task. First, VGA memory is divided into four planes that share the same address space; the 4 bits of each pixel are on different planes. So the first task is to tell the VGA which plane to use.

This is the purpose of several port I/O statements. Perversely, the VGA selects planes in different ways for reading and writing. To select the plane to be written, it uses a bit mask (0001, 0010, 0100, or 1000); to choose the plane to be read, it uses binary numbers (00, 01, 10, 11). This done, PlotSubpixel reads a byte from memory, toggles 1 bit, and writes it back. PlotDot simply calls PlotSubpixel on four neighboring points.

Given PlotDot, it's easy to construct line- and circle-drawing routines (see the text box "Bresenham's Line and Circle Algorithms" above).

Gamma Correction

One technical problem remains: How do you get a pixel to be half, or 25 percent, or 75 percent of full brightness? Brightnesses are numbered from 0 to 63, and it's naive to assume that halfway up the scale is really half brightness.

This is the familiar problem of video gamma correction. In

Listing A: *Bresenham's line and circle algorithms.*

```
{ Sign(x) is -1 if x<0, otherwise
    1 }
{ Swap(a,b) swaps values of a and
    b }

procedure Line(x,y,x2,y2:integer);
 { Note: If line is steep,
     meanings of x and y are
     swapped throughout. }
var
 i, steps, sx, sy, dx, dy, e:
     integer;
 steep: Boolean;
begin
 dx := abs(x2-x);  sx :=
     Sign(x2-x);
 dy := abs(y2-y);  sy :=
     Sign(y2-y);
 steep := (dy > dx);
 if steep then begin
   swap(x,y); swap(dy,dx);
         swap(sx,sy)
 end;
 e := 2*dy - dx;
 for i:=1 to dx do begin
   if steep then PlotDot(y,x)
         else PlotDot(x,y);
   while e >= 0 do begin
     y := y + sy;
     e := e - 2*dx
   end;
   x := x + sx;
```

```
   e := e + 2*dy
   end;
   PlotDot(x2,y2)
end;

procedure Circle(xc,yc,radius:
     integer);
var
 x,y,d: integer;
begin
 x := 0;
 y := radius;
 d := 2*(1-radius);
 while y > x do begin
   PlotDot(xc+x,yc+y);
     { also (xc-x,yc-y), etc., to
         reflect each point into
         all eight octants }
   if d + y > 0 then begin
     y := y - 1;
     d := d - 2*y + 1
   end;
   if x > d then begin
     x := x + 1;
     d := d + 2*x + 1
   end;
 end;
end;
```

agonally. If a diagonal move would land it outside the circle, it moves vertically, to get closer in. If a diagonal move would land it too close to the center, it moves horizontally to get further out.

Bresenham's key insight was how to compare these distances without doing any floating-point arithmetic. Since the equation of a circle is $y = \sqrt{(r^2 - x^2)}$, the job seems far from trivial.

But recall that to get to any point (x, y), you go through all intervening

values of x and y starting at $(0, r)$. It's easy to sum a series involving x and/or y along the way. In particular, the sum of $(1 + 2x)$ for all x and $(1-2y)$ for all y turns out to be $(x+1)^2 + (y-1)^2 - r^2$.

That's ideal, since $(x+1)^2 + (y-1)^2$ is the square of the distance from $(x+1, y-1)$—the contemplated diagonal point—to, the center of the circle, and r^2 is of course the radius. The difference between them tells you whether you're on the circle, inside it, or outside it. Because it's a difference of squares (not the same thing as a squared difference), it needs to be compared to x and y rather than to 0 or 1.

The circle algorithm in listing A is more concise than the versions usually given in the literature. It divides the circle into eight octets with a common algorithm. Studying it, I began to suspect that no one had ever hunted down all the redundancies in the arithmetic—or if they had, they kept the results strictly proprietary. There doesn't need to be even a single floating-point operation.

As shown, the algorithm computes only an eighth of a circle and obtains the rest by reflecting about the axes. To adapt the routine to devices with different horizontal and vertical pixel spacing or to draw ellipses, modify it to loop while $y>0$ rather than $y>x$, use only the first four PlotDot calls, and multiply the term 2*y by the appropriate constant when updating d.

photography, gamma is the slope of the function relating original brightness to reproduced brightness. High gamma equals high contrast. But the function is actually nonlinear, and therein lies the problem: How do you straighten it out?

Fortunately, all that's needed for the present purpose is a simple experiment. Plot a line 1 pixel wide at full brightness, a line 2 pixels wide at half brightness, and a line 4 pixels wide at one-quarter brightness. Then step back until you can't see the difference in width. If the lines look equally bright, their brightnesses are correct. Similarly, a 75 percent–bright line 4 pixels wide can be compared with a conventional line 3 pixels wide (see photo 2). That's how the values of c0, c1, c2, c3, and c4 were obtained for this program.

Gamma correction changes whenever you move to another monitor or adjust your brightness or contrast control. It also changes when images are photographed or reproduced in other ways. Fortunately, gamma correction doesn't have to be perfect for antialiasing to be effective.

What You See Is...
The best-looking two-color display from a VGA is available only through antialiasing. It is easy to implement antialiasing on the VGA because of its design. Even though antialiased displays take longer to plot, the result contains four times as much information.
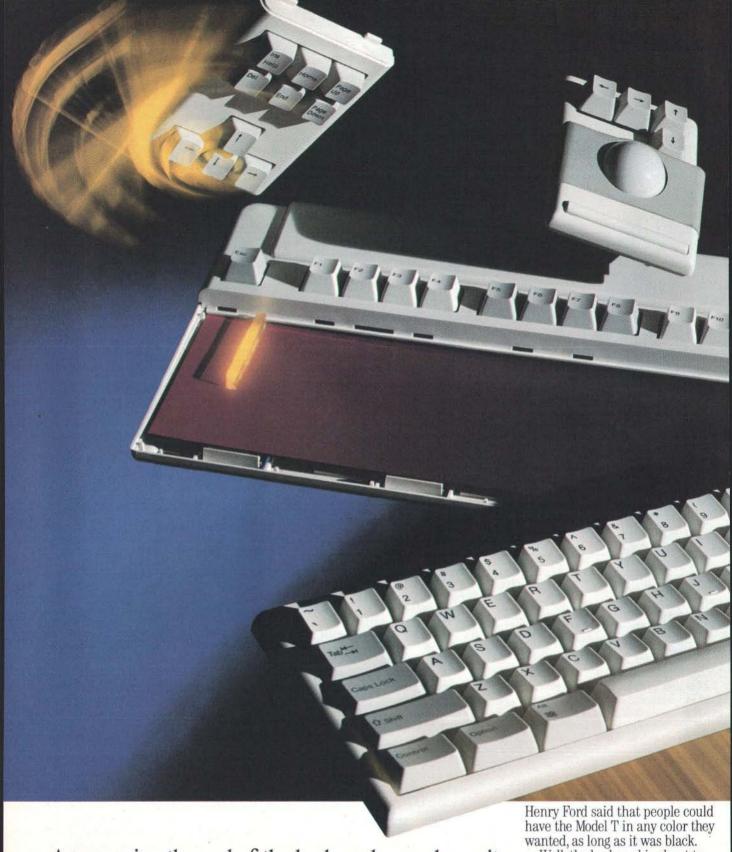
If high resolution is more important than the number of colors available, you might want to consider this approach for your application programs. ∎

Michael A. Covington does research on logic programming and natural-language processing at the University of Georgia. He is coauthor of Prolog Programming in Depth *(Scott, Foresman, 1987) and* Dictionary of Computer Terms *(Barron's, 1986) and former contributing editor of* PC Tech Journal. *He can be reached on BIX as "mcovington."*
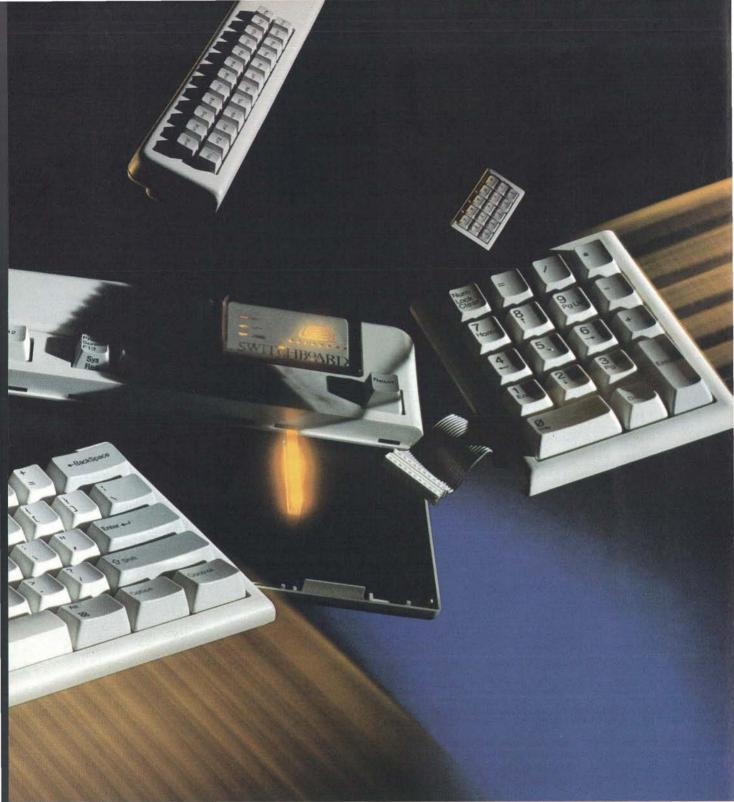
Announcing the end of the keyboard as we know it.

Henry Ford said that people could have the Model T in any color they wanted, as long as it was black.

Well, the keyboard is about to follow the Model T into history.

Because the Switchboard™ is here.

The world's first user-configurable keyboard is a radical new product based on a very simple idea: That computers should adapt to the people who use them.

Instead of the other way around.

The unique modular design lets you arrange the Switchboard's layout in any configuration that makes you comfortable.

And the more comfortable you are, the more productive you are.

To make you even more so, Datadesk is introducing a family of additional modules for the Switchboard.

Including a trackball, a vertical set of 15 function keys, a 24-key module of programmable macro keys suitable for 3270 emulation, and a dvorak alpha pad.

And that's just the beginning.

Because we're at work on still more modules designed to make the Switchboard even more powerful—and more flexible.

Speaking of flexibility, the Switchboard is the first keyboard ever to be compatible with *both* IBM® *and* Macintosh® computers.

Which makes it the only keyboard you'll ever need to know.

DATADESK

For more information or a free full-color poster, call 800/826-5398. In California 800/592-9602.

Circle 73 on Reader Service Card

OBJECT-ORIENTED PROGRAMMING FOR WINDOWS

*Using OOP to develop applications
for Microsoft Windows*

Zack Urlocker

A pplications developed for standard graphical user interfaces are becoming more and more popular, since these programs offer a consistent, easy-to-use interface, a reduced learning curve, and better integration of programs. Despite the advantages for end users, programming for a GUI with a traditional language can be overwhelmingly difficult. Object-oriented programming (OOP) reduces the complexity of GUI programming by encapsulating standard windowing behavior into predefined objects.

I'll review the basic principles of windowing environments and OOP, and then I'll walk through the design and implementation of a simple business graphics program called Chart.

GUIs: Good News for End Users and Programmers
In an increasingly competitive software market, programmers developing for GUIs have an edge over those creating traditional character-based systems. Not only is there the advantage of easy-to-use applications, but GUIs provide the programmer with a rich application programming interface (API) and, in many cases, a device-independent graphics system.

For example, programs developed for Microsoft Windows support CGA, EGA, VGA, and Hercules graphics modes, as well as hundreds of displays and printers. Since the device drivers are part of Windows, you can distribute one application that will run on any system. Writing configuration programs and device drivers, thankfully, are the GUI's responsibility, not the application's. There's also the ability to cut and paste between programs and create automatic hot links with Windows' Dynamic Data Exchange protocol. This makes it possible to integrate your application with other Windows programs you might create or with commercial applications such as Microsoft Excel.

Nevertheless, as in the laws of physics, every force has an equal and opposite force when it comes to programming for a GUI. The ease of use you supply to the end user comes at the expense of mastering the complexity of the API. Microsoft

Windows and the Macintosh Toolbox each have over 400 function calls, about 10 times as many as there are in MS-DOS. Even for the simplest of programs, you need to be acquainted with the details of standard windowing behavior, dialog boxes, menus, memory management, and event-driven programming.

Programming for a GUI with a traditional language is a little like running the Boston Marathon wearing hiking boots and a three-piece suit. Even if you're used to regular jaunts and occasional sprints, be prepared to hit the wall after about 20 miles. In programming terms, that's about 3000 lines of code, or what it might take you to create a half-decent text editor for a GUI. Even a "Hello, world!" program is about 300 lines of code when all is said and done. Of course, if you regularly undertake programming marathons, you may prefer to investigate OS/2 Presentation Manager, the triathlon of programming with over 1000 function calls.

You might be wondering if there is a way to build a higher layer on a GUI so that you don't have to know all the details of the API. OOP offers the ability to do just that, and more.

GUI Programming for the Rest of Us
OOP has always enjoyed a close relationship with graphical systems and GUIs. Smalltalk, the first completely object-oriented environment, was developed hand-in-hand with the pop-up menus, windows, and mouse that have become associated with systems like the Mac and Microsoft Windows.

Of course, the association between OOP and GUIs is by no means a historical accident. Graphical environments have an inherent complexity that makes them good candidates for the benefits of OOP. As a result, OOP has caught on widely on platforms such as the Mac, Microsoft Windows, and the NeXT machine.

Objects in a Graphical Environment
Many elements of a GUI can be conceived of as objects made up of both data and functionality. For example, a generic window
continued

Graphical environments have an inherent complexity that makes them good candidates for OOP.

includes private data to maintain its location, size, and built-in functionality so that it can be scrolled, moved, and resized. Application windows that you create should automatically include the characteristics and functionality of generic window objects.

More formally, this process is known as *inheritance*. In effect, you can create subclasses or descendants of existing kinds of objects. For example, if the Window class defines the data and functionality of generic windows, then you can create a descendant class, ChartWindow, which adds to this the ability to create and display charts. There is no need to duplicate or even understand the implementation of the generic behavior in the Window class.

Unlike traditional function libraries, where you are locked into the limitations imposed by the original programmer, in an object-oriented system you're free to override any of the inherited functionality. This gives you the ability to further customize objects for more specialized behavior. For example, if your application requires that the main window cannot be resized or closed, you can simply override these capabilities.

Many object-oriented languages have numerous predefined classes of objects that give the programmer a head start in developing applications for GUIs. Actor, for example, includes classes of windows, dialog boxes, and graphics, as well as data management. The standard classes provide the generic capabilities common to all Windows applications, thereby eliminating the need to become involved in the details of memory management, pointers, handles, and most of the API.

Beyond Windows and Dialog Boxes

If OOP only allowed you to simplify the development of user interfaces for GUIs, it would still provide a dramatic increase in programming productivity. However, OOP goes beyond windows and dialog boxes. In an OOP system, all the logical entities that make up your application can be created as objects. Since objects are both data and functionality, they more closely resemble the real-world entities you're dealing with, giving you a solution that is easier to understand.

For example, an accounting program would be made up of objects that correspond closely to the real world, such as ledgers, transactions, and accounts. In the Chart program, you will create pie chart and bar chart objects, as well as a chart window and chart dialog box. These objects embody all the functionality that you expect of a chart, window, and dialog box, respectively.

Objects can be made up of other objects declared as instance variables. Instance variables correspond to fields in a data structure. For example, a Chart object has instance variables called data and labels that are arrays of integer objects and string objects, respectively.

Handling Window Events and Messages

Most graphical environments are event-driven systems. For example, when the user presses a key or clicks the mouse, this is an event that is trapped by the GUI and handled by the application. On the Macintosh, for example, applications contain a main loop that polls for events and determines how to handle them.

In Microsoft Windows, events are more generalized and are

MICROSOFT WINDOWS MESSAGES

Table 1: Windows messages and their associated events and parameters.

Message	Event	wParam	lParam
WM_KEYDOWN	Key press	Key code	Repeat count
WM_LBUTTONDOWN	Left mouse button	State	Location point
WM_SIZE	Window resize	Type	Width/height
WM_ACTIVATE	Window selected	Type	Handle
WM_HSCROLL	Window scrolls	Code	Position
WM_CLOSE	Window closed	—	—
WM_PAINT	Repaint the window	—	Paint struct
WM_COMMAND	Menu command	Item	Type

USER EVENTS IN MICROSOFT WINDOWS

User clicks the mouse button in a window.

Window sends WM_LBUTTONDOWN message with location point as a parameter.

The ChartWindow object gets the WM_LBUTTONDOWN message and responds appropriately.

Figure 1: *Most events are initiated by the user. In the example, the user presses the mouse button, causing a* WM_LBUTTONDOWN *message to be sent to the Window in which the user clicked.*

treated as messages. When an event occurs, Microsoft Windows sends a message to the appropriate window. Figure 1 shows how events are processed in Microsoft Windows.

Microsoft's window messages are sent with two parameters to supply additional information. These are known as the word parameter or wParam, a 2-byte number, and the long parameter or lParam, a 4-byte pointer. For example, when you press a key, a WM_KEYDOWN message is sent to the window, with the wParam giving the value of the character. Table 1 shows examples of other Windows messages and the meaning of the parameters.

When developing a Windows application in C, you process the messages by writing a function known as the WndProc, which contains a switch statement to determine what the message is and how it should be handled. An example of the WndProc function is shown in listing 1.

Messages and Methods

In an object-oriented system, all computation is done by sending messages to objects. Like messages in a windowing environment, a message is simply a request to an object to perform some action. To draw a chart, for example, you send a draw message to a Chart object:

```
draw(myChart, context); /* Actor code */
```

In this example, myChart is the receiver of the draw message, and context is an argument indicating the display context where the drawing is to take place. To respond to the draw message, the Chart object must define a draw method. Thus, a message corresponds to a function call, and a method corresponds to a function definition in a procedural programming language.

One major distinction between a message and a function call in a procedural language is that the different objects can respond to the same message in different ways. For example, all graphical objects, including rectangles, ellipses, points, and charts, can respond to the same general-purpose draw message. In a procedural language such as C, you would have to know what type of data you are dealing with and then call the appropriate function. In an object-oriented language, the system handles this task for you.

The ability of objects to respond to general-purpose messages in different ways is known as *polymorphism*, meaning "many behaviors." Polymorphism lets you write more generic, reusable code, since the objects themselves will handle the details.

Since object-oriented languages already handle messages, Windows messages can be handled easily. In Actor, Windows messages are automatically trapped and sent as Actor messages to the appropriate window object, with no code required by the programmer. Higher-level Actor methods hide many details of processing the Windows messages.

For example, the Window class defines a WM_PAINT method to trap the message requesting the window to repaint itself. The WM_PAINT method will allocate memory for repainting, and then it will send an Actor paint message with a display context in which to draw. Code for the WM_PAINT method is shown in listing 2.

Since this code is already built into the system, descendant classes merely have to respond to an Actor paint message. They do not have to directly allocate or deallocate memory for the display context.

In order to redraw the chart contained in a ChartWindow, the paint method is written as shown in listing 3. Of course, this

Listing 1: WndProc *for Microsoft Windows in C. This code shows the logic of how Windows messages are processed in C. Processing details are in the user-defined functions* newfile(), openfile(), savefile(), help(), hittest(), *and* keypress().

```
long FAR PASCAL WndProc(hWnd, message, wParam, lParam)
HWND hWnd;                        /* window handle    */
unsigned message;                 /* type of message */
WORD wParam;                      /* word parameter  */
LONG lParam;                      /* long parameter  */
short x, y;                       /* mouse location  */

{
    FARPROC lpProcAbout;          /* ptr to About function */
    HMENU hMenu;                  /* handle to system menu */

    switch(message)               /* what event has occurred? */
    {
        case WM_SYSCOMMAND:       /* system menu commands  */
            if (wParam == IDABOUT) /* About... menu command */
            {                     /* display a dialog box  */
                lpProcAbout = MakeProcInstance(About, hInst);
                DialogBox(hInst, "AboutBox", hWnd, lpProcAbout);
                freeProcInstance(lpProcAbout);
                break;
            }
            else                  /* do default processing */
                return(DefWindowProc(hWnd, message, wParam, lParam));

        case WM_COMMAND:          /* menu commands */
            switch(wParam)        /* what menu item? */
            {
                case CW_FILE_NEW: /* various menu commands */
                    newfile();
                    break;
                case CW_FILE_OPEN:
                    openfile();
                    break;
                case CW_FILE_SAVE:
                    savefile();
                    break;
                case CW_HELP:
                    help();
                    break;
            }
            break;

        case WM_LBUTTONDOWN:      /* left mouse button */
            x = LOWORD(lParam);   /* x coordinate      */
            y = HIWORD(lParam);   /* y coordinate      */
            hittest(hWnd, x, y);  /* hit test logic    */
            break;

        case WM_KEYDOWN:          /* key press         */
            keypress(hWnd, wParam); /* key press logic  */
            break;

/*
 * other messages go here...
 */
        case WM_DESTROY:          /* quit the application */
            PostQuitMessage(0);
            break();

        default:                  /* default processing */
            return(DefWindowProc(hWnd, message, wParam, lParam));
    }
    return(NULL);
}
```

code simply sends a draw message to the Chart object, which must in turn handle the details of drawing.

Objects as Abstract Data Types

By defining a complete set of messages that objects understand, you can shield the physical representation from users. This makes it possible to treat objects as abstract data types that can be manipulated without regard for their physical implementation.

Listing 2: *Trapping Windows messages in Actor. This method defines default processing of the WM_PAINT message in the* Window *class. It allocates memory for the paint structure, creates a display context, and sends an Actor paint message to the window. When finished, the memory used for painting is deallocated. The* self *parameter refers to the window that was clicked in;* dc, hPS, *and* lpPS *are local variables referring to the display context, handle to the paint struct, and long pointer to the paint struct, respectively. If you define any descendant classes, they should define a paint method to handle redrawing in the display context.*

```
Def  WM_PAINT(self, wParam, lParam | dc, hPS, lpPS)
{
  hPS := asHandle(paintStruct);        /* get a paint struct  */
  lpPS := globalLock(hPS);             /* lock down memory    */
  dc := Call BeginPaint(hWnd, lpPS);   /* get display context */
  paint(self, dc);                     /* send Actor message  */
  Call EndPaint(hWnd, lpPS);           /* finish painting     */
  globalUnlock(hPS);                   /* unlock memory       */
}
```

Listing 3: *The paint message for* ChartWindow, *which is used to redraw the chart when necessary.*

```
Def paint(self, context)
{
    if chart then
       draw(chart, context);   /* tell the chart to draw */
    endif;
}
```

Listing 4: *The* init( ) *function initializes a new* Chart *object and its instance variables. The data and labels collections have an initial size of 10 but will automatically grow to accommodate additional data. The variables are used as follows:* data *refers to a collection of data values;* labels *is a collection of text labels;* space *is an integer indicating the spacing between items;* scale *is a point indicating the scale to be used; and* lead *is a point indicating the lead spacing to be used. The* resetScale( ) *function sets the scale instance variable to a point. Descendant classes may override this method. The* resetLead( ) *function sets the lead instance variable to a point. Descendant classes may override this method.*

```
Def init(self)
{
  data := new(OrderedCollection, 10);
  labels := new(OrderedCollection, 10);
  space := 10;
  resetScale(self);
  resetLead(self);
}

Def resetScale(self)
{
  scale := 10@10;
}

Def resetLead(self)
{
  lead := 10@10;
}
```

Listing 5: *Code for* reScale *and* draw *methods of* Charts. *The first function,* reScale( ), *rescales the chart based on data and area. Descendant classes will define the* adjustScale *method. The second function,* draw( ), *is used to draw a chart in the given display context by drawing the data and the labels.*

```
Def reScale(self | max)
{
  max := 0;                    /* find the maximum */
  do(data,
     {using(item)
      if item > max
            then max := item;
      endif;
   });
  if max > 0 then              /* rescale according to max */
    resetLead(self);
    adjustScale(self, max);
  endif;
}

Def draw(self, context)
{
  if area then
    drawData(self, context);
    drawLabels(self, context);
  else
    error(self, stackTop( ), #areaNilError);
  endif;
}
```

Listing 6: *In the* HBarChart *class: This* resetLead( ) *overrides the inherited* resetLead *method to allow more room for the labels on the left edge of the chart. The* adjustScale( ) *function adjusts the scale based on the data and area. The* drawLabel( ) *function draws the labels, computes the positions, and converts the label into a string. The* drawData( ) *function draws the data elements in the given display context. It computes the positions based on scale and lead (x, y, and* dRect *are temporary variables).*

```
Def resetLead(self)
{
  lead := 90@20;
}

Def adjustScale(self, max)
{
  scale := point((x(area) - 2 * x(lead)) / max, 15);
}

Def drawLabel(self, context | str, x, y)
{
  x := 1;                             /* left edge */
  do(size(labels),                    /* for each label */
  {using(i)
    y := y(lead) + i * (y(scale)+space); /* vert. location */
    str := asString(labels[i]);       /* convert it */
    drawText(self, x, y, str, context); /* draw it */
  });
}

Def drawData(self, context | x, y, dRect)
{
  x := x(lead);                       /* left edge of bar  */
  do(size(data),                      /* for each item     */
  {using(i
    y := y(lead) + i * (y(scale)+space); /* vert. location */
    dRect := rect(x, y,               /* create a rect     */
                  x + data[i]*x(scale),
                  y + y(scale));
    fill(dRect, stock(GRAY_BRUSH), context); /* draw it */
  });
}
```

For example, the physical representation of Chart objects is considered private. Users of Chart objects do not need to know how Charts are implemented or what their instance variables are. Instead, they send messages to Chart objects to perform certain actions, set their private data, or return a value. These messages define the "public protocol." In the case of Chart objects, the protocol might include messages such as draw, size, and reScale. You can then send messages to objects telling them what to do without getting caught up in the details of how they do it. Figure 2 shows a diagram of a Chart as an abstract data type.

The advantage of separating the public protocol from the private implementation is that it reduces the dependencies in the system to a well-defined interface. You can easily change the internal implementation of Chart objects knowing that users of Charts will not be affected.

Designing the Chart Application

Since the Chart application is quite simple, it has only a few different classes, including ChartWindow, ChartDialog, and several kinds of charts that descend from the Chart class.

Figure 3 shows a class tree for the Chart application. Table 2 describes the classes. The Chart class is used to define all the common data and functionality among the various kinds of charts. For example, all charts should respond to the same messages: draw, reScale, print, load, save, and so on.

Similarly, all charts contain the same instance variables: data and labels maintain the numeric data and the string labels that will be plotted; area and scale are point objects; and lead and space are integer objects.

The only real difference between the different kinds of charts is how they implement the details of drawing the data and labels. All the common code is stored in the Chart class and automatically inherited by the descendant classes. Since there is nothing chart-like already in the system, the Chart class descends directly from the Object class, inheriting only minimal object functionality.

The code to initialize new Chart objects is shown in listing 4. Initialization is broken down into three methods to allow descendant classes to override parts of the initialization as necessary. For example, horizontal bar charts may have a need for different lead spacing to accommodate text labels. Thus, they can override the resetLead method inherited from the Chart class with their own methods. Note that the inherited code of class Chart exists only once in the application; it is not copied into the other classes. Inheritance lets you reduce the size of the application by factoring common parts into ancestor classes.

A chart will receive messages to draw or reScale from the ChartWindow in which it is contained. For example, if the user resizes the ChartWindow, the ChartWindow will send a reScale message to the chart. Listing 5 shows the code for the draw and reScale methods using two interesting messages, do and error.

The do message lets you loop over a collection's items and perform a block of code without having to keep a count. For example, the code

```
do(aCollection, /* iterate over a collection */
{using(item) /* item is the loop variable */
printLine(item) /* print each item */
});
```

iterates through all items in a collection, printing each. The do message is understood by all the system's collections—Arrays,
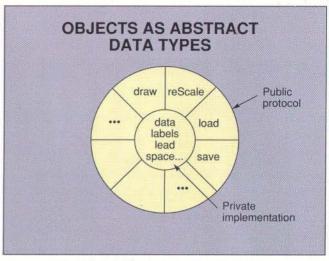
Figure 2: *The* Chart *object has a private internal implementation and a public protocol layer.*
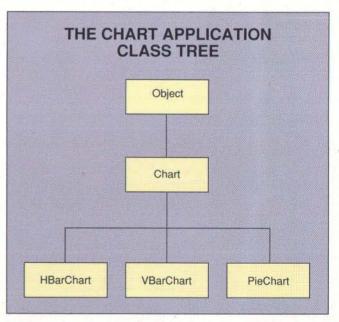
Figure 3: *The* Chart *class defines common data and functionality among the descendant classes.*

CLASSES IN THE CHART APPLICATION

Table 2: A description of each class used in Chart, the sample application.

Class	Purpose
ChartWindow	Display the chart, handle user interaction.
ChartDialog	Allow the user to edit the chart data points.
Chart	Formal class to define common chart functions.
HBarChart	Horizontal bar chart.
VBarChart	Vertical bar chart.
PieChart	Pie chart.

Listing 7: *The Chart application header file,* chart.h, *contains the constants for the charting program.*

```
#define TWO_PI          6.283185307   /* for pie charts */

#define CW_FILE_NEW     501           /* menu item IDs */
#define CW_FILE_OPEN    502
#define CW_FILE_SAVE    503
#define CW_HELP         600

#define CW_HELP_BOX     801           /* dialog box IDs */
```

Listing 8: *Menu handling in the* ChartWindow *class. This* init( ) *automatically initializes new* ChartWindow *objects by setting up the menus and their logic. The menu definitions are loaded from resources. The actions dictionary provides the menu logic. Keys in the dictionary give the menu item ID. Values in the dictionary give the message to be performed. The* command( ) *function handles menu commands using a data-driven approach. The first argument,* menuItem, *indicates the menu item ID. It looks up* menuItem *in the actions table, and if it is found, performs the appropriate action. It ignores* lParam.

```
Def init(self)
{
    chart := new(VBarChart);       /* default chart type */
    addAbout(self);                /* add an "about box" */
    loadMenu(self, "CWMenus");     /* load menu resources */
    setMenu(self, hMenu);

    actions := new(Dictionary, 4); /* add menu item logic */
    add(actions, CW_FILE_NEW, #newChart);
    add(actions, CW_FILE_OPEN, #restoreFile);
    add(actions, CW_FILE_SAVE, #saveFile);
    add(actions, CW_HELP, #help);
}

Def command(self, menuItem, lParam)
{
    if actions[menuItem] then
        perform(self, actions[menuItem]);
    else
        beep( );
        errorBox("Error", "Undefined command: "
            + asString(menuItem)));
    endif;
}
```

OrderedCollections, SortedCollections, and others. The error message lets you define your own error conditions. You can flag different error conditions and easily define your own error handlers for different kinds of objects.

Chart's Descendant Classes

To fully implement the charts, each descendant class must define the drawData, drawLabels, and adjustScale methods as shown in listing 6 for the HBarChart class. These methods implement the rest of the protocol for Chart descendant classes. The drawData method calculates the location of each bar in the chart as a rectangle from the left edge. Then the rectangle is drawn in gray by sending it a fill message.

The code for VBarChart is similar to that for HBarChart. For the PieChart class, the implementations of drawLabels and drawData are different, since drawing a pie chart can be done by calling a Microsoft Windows function directly. The drawData method for PieChart encapsulates the direct call to the Windows Pie function as shown below.

continued

```
/* call the Windows Pie function directly */

Call Pie(context, xLeft, yTop, xRight, yBottom
    xStart, yStart, xEnd, yEnd);
```

The ChartWindow Class
The ChartWindow, of course, inherits most of its functionality from the Window class. In addition, it responds to the paint message, as shown previously in listing 3.

The ChartWindow provides the user interface to the Chart application. This includes menu selection, resizing the window, and clicking the mouse to edit data points. Whenever a menu event occurs, Windows will send a WM_COMMAND message, with the wParam indicating the menu item that was selected. In C programs, normally a case statement is used to test the value of the wParam and take the appropriate action.

In Actor, a more object-oriented approach can be used. Typically, a table is used to look up the method to be performed based on the menu item selected. If the user selects, say, the File New command, the wParam will have the value defined by the constant CW_FILE_NEW, and the newChart method should be performed. The menu item constants are defined in a header file, shown in listing 7, which uses C-style #define statements. The numbers identified with the menu items are arbitrary, but they must be unique.

Listing 8 shows the code used to initialize the ChartWindow's instance variables, including the chart and the actions lookup table. It also loads the resources that define the look of the menu. The command method looks up the wParam in the actions table and, if it finds it, performs the method.

The Main Program
Since the ChartWindow is really the centerpoint of the application, all the main program really has to do is create a new ChartWindow object and show it on the screen. This requires only two lines of code:

```
CW := defaultNew(ChartWindow, "Chart)";
show(CW, 1);
```

Although I haven't shown all the Chart code, I've highlighted the most interesting aspects of an object-oriented approach to GUI programming—encapsulation, inheritance, message sending, and polymorphism.

An object-oriented approach can help reduce the complexity of GUI programming and can help give you a competitive edge in the marketplace. OOP increases your productivity by providing predefined classes that encapsulate the API. By using inheritance, you can customize the existing classes and extend them for your own applications. In addition, since GUIs are event-driven, they map well to the message basis of object-oriented systems. ∎

Editor's note: *The Chart source code is available in a variety of formats for those who would like to explore it further (see page 5 for details). Chart is written in Actor, which is an OOP language for Microsoft Windows. The* Chart *classes can be used as they are, or descendant classes can be created for your particular needs.*

Zack Urlocker is manager of developer relations at the Whitewater Group, the creators of Actor. He holds a master's degree in computer science from the University of Waterloo, where his research was in the area of programming environments. He can be contacted on BIX c/o "editors."

There's gold

Now, Quarterdeck's new

Memory is gold.

And like gold, some of it is hidden away inside your computer. For years, we've been working toward putting it all under your control. And now we can.

Now you can make today's more powerful programs run without giving up network and mouse drivers and TSRs.

Introducing Manifest—the Quarterdeck memory analyzer

Many PC users know there are nuggets of memory sitting unused in most PCs. But those little pieces of memory can add up to 130K!

That's why Quarterdeck Office Systems, publisher of DESQview, developed a new utility that helps you find and use this memory. It's called Manifest. And it does for memory what PC Tools does for disks. For under $60.

Quarterdeck's seven years of memory expertise made Manifest

Manifest guides you deep inside your PC.

It locates unused (or underused) memory and suggests where you could load networks, buffers, mouse drivers, TSRs and other utilities to increase performance. It even analyzes what type and amount of RAM you have available, and which portions of your memory are faster.

> Administering a number of PCs? Manifest's diagnostic and reporting capabilities reduce technical support time . It not only identifies problems but helps to solve them.

Manifest shows you how your memory works. Here's the first megabyte of RAM, showing unused areas.

Your current memory is full of holes. Our tools can fill blocks of unused addresses between 640K and 1024K to free up memory your programs can use.

Manifest shows you the contents of AUTOEXEC.BAT and CONFIG.SYS files. That can be a big help when diagnosing problems. Manifest tells you all about your hardware, too —from your cpu type to what boards you have installed. Manifest even tests memory speed.

And it runs benchmark tests on expanded memory boards so you can make informed buying decisions.

You won't need a PhD to understand what you're doing. Manifest has an interactive 'manual' that tells you how to use the program and what benefits you'll get.

And unlike a lot of hot new software, Manifest works on virtually any PC: 8088, 8086, 80286 or 80386. It's a productivity breakthrough from the memory experts at Quarterdeck.

Introducing QRAM—the Quarterdeck memory optimizer

End RAM cram in your 8088, 8086 or 80286 PC once and for all. QRAM (pronounced cram), is a package of utilities that gives you unprecedented control over memory, letting you set up your

memory the way it will work best for you.

If you have EMS 4.0 or EEMS boards, QRAM can find unused addresses and 'map' memory to those addresses. Then it looks at your AUTO-EXEC.BAT and CONFIG.SYS files and figures out what TSRs, network and mouse drivers and DOS resources can be loaded high and where.

And, like all Quarterdeck memory products, QRAM is compatible with the Microsoft XMS specification used by Windows 286, V. 2.x.

If your PC has 'shadow RAM,' there's even

QRAM optimizes your memory performance by moving utilities and drivers out of the area between 0K and 640K — freeing it up for your programs to use.

more gold in your PC. QRAM finds the unused parts and puts them under your control.

And if you have an EGA or VGA-equipped PC and don't need graphics at the moment, QRAM will make an additional 96K 'nugget' of memory available! When you need graphics again, QRAM will switch you back to graphics mode! Think how helpful that will be for those big dBASE files.

It can't work miracles, but if there's memory available anywhere, QRAM lets you use it to increase your PCs speed and performance.

QRAM is available bundled with Manifest for just a few dollars more than Manifest alone.

Manifest and QRAM—two more examples of Quarterdeck's commitment to mining the most productivity out of the PC and software you own today.

REELING IN THE DATA

Tape backup systems are no longer just for the wealthy or the overanxious

O ne day it will surely happen: Your system's hard disk drive will die of old age, mysteriously lose data, or succumb to power glitches or (worse yet) a fire. And, if you're like the majority of microcomputer users, you'll lose precious data because you didn't make a recent backup.
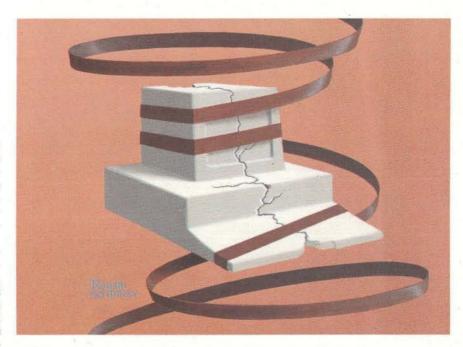
This situation is, of course, preventable if you're willing to spend a few minutes each day backing up your data, either to floppy disks (using a high-speed floppy disk backup utility) or to tape. In this installment of Under the Hood, I'll focus on the latter medium and describe how the most popular types of tape backup units work.

Why Tape?

Magnetic tape is ideal for situations in which you need to store large amounts of data without changing the medium—typically, unattended backup or the distribution of large programs or databases. The lowest-capacity tape drive you're likely to find will store 40 megabytes of data on a single cartridge. High-end helical-scan units from Exabyte store 5 gigabytes per cartridge and will probably increase to 8 gigabytes before the end of the year.

By contrast, a standard AT-type 5¼-inch floppy disk can hold only 1.2 MB of data; its 3½-inch counterpart can hold only 1.44 MB. Backing up a large hard disk drive with a utility like Fastback from Fifth Generation Systems can be a true test of one's disk-swapping patience.

Extra-high-density floppy disks—like those from Insite Peripherals and Brier Technology—are available, and you may find them useful as a backup medium.

However, these disks will hold at most 40 MB, the same as the least capacious tape, and aren't widely available yet. Tape drives, on the other hand, are available from numerous manufacturers, and media for them are available at virtually all computer stores. And tape is generally less expensive, megabyte for megabyte, than floppy disks.

One area where tape is *not* superior to floppy disks is durability. Floppy disks, which have both a thicker substrate and a heavier magnetic coating, have a better chance of coming through extremes of temperature (e.g., fire) intact.

What about WORM (write once, read many times) drives and read/write optical disk drives? So far, the main disadvantage of these drives is cost: $4000 to $8000 per unit, plus $50 or more per disk. The advantage of such drives is that they're useful for other kinds of storage between backups.

Disks of any kind—magnetic or optical—will always be far superior to tape when random access is required. It's much easier and faster to step from track to track on a disk than to fast-forward and rewind a tape drive. But in backup applications, which primarily involve the transfer of long, continuous streams of data, tape truly comes into its own.

Varieties of Tape Formats

The four tape formats most often used for data storage are half-inch reel-to-reel (most often used on minicomputers and mainframes), quarter-inch cartridge, 8-mm cartridge, and 4-mm DAT (digital audio tape).

Of these four, half-inch reel-to-reel (often called nine-track) is the most mature technology. As the name suggests, data is laid down in nine parallel tracks at 800, 1600, 3200, or 6250 bits per inch. The tracks are read in parallel using a head with nine distinct elements, making data transfers extremely fast, even at relatively low tape speeds. A full

INSIDE A QUARTER-INCH TAPE CARTRIDGE

Belt guide rollers (2)

Supply hub   Tape shown in   Flat drive belt   Take-up hub
              beginning-of-tape
              position

Read/write
head

Guide pins (3)

Belt capstan

Rubber-covered
drive roller

Drive motor   Direction of forward   Light sensing   Fixed tape
               tape motion   guides (2)

Mirror

Base plate

Tape path
(oxide out)

File
protect
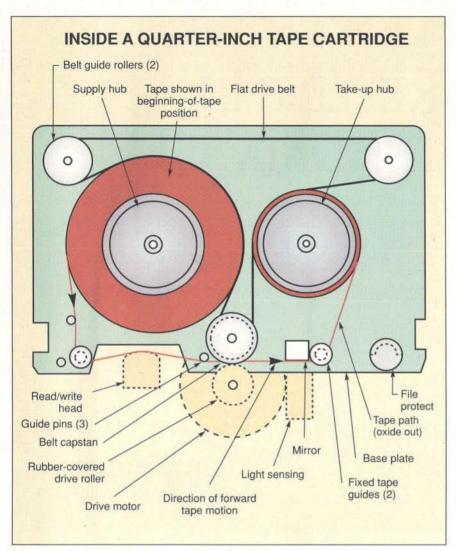
Figure 1: *Quarter-inch tape cartridges use a clever drive mechanism that moves the tape directly, eliminating the need for separate drives and clutches for the hubs.*

reel of tape recorded at the highest density typically holds about 145 MB.

Nine-track tape is time-tested and almost 100 percent interchangeable between virtually all the drives that handle it. It's frequently used to transport data between machines as well as to back up disks. But nine-track drives—especially good ones—aren't cheap, and the big reels of tape are awkward and not likely to be available at your local computer store. Also, while they offered more capacity than cartridge units until a few years ago, they're now falling behind. Perhaps this is the reason you don't see nine-track tape on many personal computers nowadays.

Clever Cartridges

The most common medium used by tape backup units is quarter-inch tape, which

is enclosed in one of two sizes of cartridges: a DC2000-class minicartridge, which is about 2½ by 3 inches, or a DC600-class cartridge, which is the same shape but about 30 percent bigger in each dimension. (A new class of cartridge, DC9135, will look like a DC600 but will have 900-oersted tape, which can handle denser recording schemes, inside.) The exact specifications for the cartridges, originally defined by 3M, are now embodied in ANSI standards.

Each cartridge (see figure 1) contains two tape hubs, a flexible drive belt that passes partway around each hub, tape guides, and a small mirror assembly (used by the drive to detect holes at the beginning and end of each tape). A small door (not shown) protects the cavity where the read/write head enters the cartridge from dust and debris.

The mechanism inside a quarter-inch cartridge is clever and unique. Even though it's been around for more than a decade, you won't see anything like it in an audiocassette or a videocassette, probably for patent reasons. Note that the drive belt actually contacts the tape and moves with it partway around each hub, ensuring that the linear velocity of the tape is equal to that of the belt no matter how much tape is on each hub. The angular velocities of the hubs, which aren't directly driven, vary as necessary. This makes the drive extremely simple at the expense of a small amount of additional hardware in the cartridge.

All the drive needs to do to position the tape is rotate a pulley called the *belt capstan*; the pulley moves the belt, which in turn moves the tape. There's no need for a tape capstan, a pinch roller, or independent drive mechanisms and clutches for the hubs. (An ordinary Philips-type audiocassette deck, by contrast, must have all these things.)

Cartridges within a class may have different lengths. Many users have been pleasantly surprised to find extra-long tapes appearing on computer store shelves, letting them increase the capacity of an existing drive by up to 50 percent without buying new hardware. A QIC-40 drive (nominal capacity of 40 MB) can often record 60 MB on such tapes, and a QIC-80 drive (80 MB) can get 120 MB on one tape.

QIC and Easy

Standards that describe how to record on quarter-inch cartridges are published by the organization Quarter Inch Cartridge Drive Standards (311 East Cabrillo St., Santa Barbara, CA 93101). So far, the organization (QIC for short) has adopted 36 standards, specifying such things as interfaces between computers and tape drives, tape formats, recording-head properties, error-correction codes, data-compression algorithms (see the text box "Data Compression Doubles Tape Capacity" at right), and SCSI command sets for tape drives (see the table for a partial list; if you'd like a complete list, contact QIC).

Some of these standards have become pervasive enough that makers of drives for other tape formats have adopted them. For instance, most 4-mm DAT drives used for tape backup implement the set of SCSI commands specified in the QIC-104 standard and will probably move to QIC-121 (a SCSI-2 command set) in the future.

As a result of this standardization

Data Compression Doubles Tape Capacity

Data compression is a highly effective way of increasing the capacity of a tape drive. Unfortunately, most software algorithms with good compression ratios can process at most 64K bytes per second—and would therefore constitute a major bottleneck in a tape backup system.

To solve this problem, a company called Stac Electronics (Pasadena, CA) has developed a data-compression/decompression chip with an average throughput of 750K bytes per second and an average compression ratio of 2-to-1. The chip, which uses a modified Ziv-Lempel algorithm, achieves this ratio using only 16K bytes of RAM, only 2K bytes of which is used to hold the string table for the compression algorithm.

Even though the company refuses to divulge exactly how the proprietary algorithm works, the organization Quarter Inch Cartridge Drive Standards (Santa Barbara, CA) has adopted it as the QIC-122 standard. Should vendors wish to implement the algorithm in software only, rather than buying a $50 chip to put in every drive, Stac Electronics will license object code for the cool sum of $25,000.

Because other manufacturers may wish to go their own way rather than embrace one manufacturer's proprietary solution, the QIC committee has established a registry of "data algorithm identifiers" that can be written on a tape to show how the data is encoded. A tape backup system can use this code to apply the correct algorithm—if it knows it—to the data on the tape while reading it.

SOME KEY QIC STANDARDS

A partial list of the 36 standards adopted by Quarter Inch Cartridge Drive Standards.

Interfaces

QIC-02	Quarter-inch cartridge tape drive intelligent interface
QIC-36	Quarter-inch cartridge tape drive basic interface
QIC-104	Implementation of SCSI for QIC-compatible sequential storage devices
QIC-121	Implementation of SCSI-2 for QIC-compatible sequential storage devices

Commonly used tape formats

QIC-24	Serial-recorded magnetic tape cartridge for information interchange (60 MB)
QIC-40	Flexible-disk-drive-controller-compatible recording format for information interchange (40 MB)
QIC-80	Flexible-disk-drive-controller-compatible recording format for information interchange (80 MB)
QIC-120	Serial-recorded magnetic tape cartridge for information interchange (125 MB)
QIC-150	Serial-recorded magnetic tape cartridge for information interchange (150 MB)
QIC-525	Serial-recorded magnetic tape cartridge for information interchange (525 MB)
QIC-1350	Serial-recorded magnetic tape cartridge for information interchange (1.35 GB)

Data compression

QIC-122	Data-compression format for quarter-inch data cartridge tape drives
QIC-123	Registry of data algorithm identifiers for quarter-inch cartridge tape drives

A MESSAGE TO OUR SUBSCRIBERS

Figure 2: *Slanted tracks maximize use of the tape's surface area; alternate azimuth angles minimize cross talk.*

effort, tapes recorded on one manufacturer's drive can be read in another, provided that the industry-standard format is used. (Almost all manufacturers have their own enhanced, proprietary, and incompatible formats in addition to the standard ones.) In December 1989, an independent laboratory called Pericomp (Natick, MA) certified QIC-40 cartridge drives from three manufacturers—Archive, Mountain Computer, and Colorado Memory Systems—as fully compatible. Other manufacturers (such as Alloy Computer Products and Everex Systems) are expected to follow, and Pericomp has begun to certify QIC-80 drives as well.

"Floppy Tape"

The QIC-40 and QIC-80 standards are sometimes called the "floppy tape" standards, because they're designed to use the computer's existing floppy disk drive controller to read and write tapes. This is less expensive than providing a custom controller and may save backplane slots, which is especially important for users with "baby" motherboards. However, this approach has some drawbacks.

First, since the maximum bit clock rate of many floppy disk drive controllers

is either 250,000 bps (as on the IBM XT) or 500,000 bps (as on the IBM AT), there's an upper bound on how fast data can be stored and retrieved without a special controller.

Second, since many computers (e.g., the AT) were designed to have at most two floppy disk drives, they don't provide select lines for any more devices. So, in systems that have two floppy disk drives, clever software and hardware tricks are required to let the tape drive share the interface. One commonly used technique is to tell the floppy disk drive controller to select *no* drive; circuitry on the drive or an interface card detects this condition and activates the tape. Pulses on the floppy disk drive controller's step line are often used to send commands to the tape drive; the drive counts the pulses to determine the command.

If, like me, you have an IBM AT compatible that has high- and low-density floppy disk drives and no more cutouts in the chassis, you may suffer from another problem when buying one of these normally inexpensive tape backup units: sticker shock. An external drive can cost as much as $400 more than an internal drive due to the costs of a case, power supply, and FCC-certified cable.

Still, if you find yourself in this situation, resist the temptation to write (and potentially ruin) 360K-byte floppy disks in the high-density drive. Instead, spend the extra money or consider moving to a high-density 3½-inch drive, which will not have compatibility problems with low-density 3½-inch disks.

In the QIC-40 and QIC-80 standards, the floppy disk drive controller sees blocks of data on the tape as if they were disk sectors 1K byte in length. The standards specify how the tracks of the tape are mapped to the sectors, tracks, and sides of the imaginary floppy disk. The tracks of the tape are read and written sequentially in a serpentine fashion. The cyclic redundancy checks that are generated by the floppy disk drive controller are used for error detection, and a Reed-Solomon code is used to correct errors where possible.

Higher-numbered QIC formats move away from using the floppy disk drive controller to read and write and are usually implemented within SCSI-based tape drives. This results in full SCSI speed and greater interchangeability of drives—but not necessarily tapes—among vendors.

Triple Helix

At this writing, the only available backup systems that can store more than a gigabyte per cartridge use helical-scan technology (i.e., the same kind that is used in VCRs and DAT recorders) to achieve this and higher capacities.

Figure 2 shows how a helical-scan drive reads and writes data. The tape wraps around a spinning drum whose axis is tilted at a small angle (typically 5 to 6 degrees) to the tape. The heads, which are mounted on the drum, traverse the tape in long, diagonal stripes (actually, sections of a helix). Each section is about nine times the width of the tape in length.

This is the point where the 8-mm and 4-mm designs diverge, however. Exabyte's 8-mm helical-scan technology, which is based on an 8-mm analog VCR mechanism supplied by Sony, uses three heads (a servo head, a read head, and a read-after-write head) on the drum, plus a separate erase head that erases the full width of the tape at once.

Tracks are written individually, and they contain 8K bytes of data each. The tape wraps more than halfway around the

head—211 degrees, to be exact. The current storage capacity of the Exabyte drive is 2.3 gigabytes per cartridge, with an increase to 5 gigabytes expected sometime this year.

The DDS (Digital Data Storage) format, developed and licensed by Hewlett-Packard and Sony, currently appears to be the dominant data format for 4-mm DAT. DDS uses the innards of a DAT deck, with four heads on the drum: two write and two read-after-write. Tracks are written in pairs (called *frames*) that actually overlap slightly.

The heads can distinguish the tracks, however, because they are recorded at different *azimuth angles* (i.e, with heads that are tilted at different angles relative to the direction of the track). Each frame contains 8K bytes, and the tape wraps only 90 degrees around the drum, which the designers claim helps to reduce tape wear. DDS tapes currently hold 1.3 gigabytes each.

Another format for DAT, called Data/DAT, is now under development by another consortium. While it is similar to DDS in many ways, it adds one important *continued*

Circle 291 on Reader Service Card

Create, Capture, & Present Your Imagination!

. . . with sonic digitizers.

You can create the simplest, contour images or develop complex, animation data bases. Capture graphics from any medium up to 60" × 72", or trace three dimensional objects or physical movement within a 9 foot × 9 foot × 9 foot cube. What you choose to input is only limited by space...your output is only limited by your imagination.

With a sonic digitizer from SAC, you can input data directly from any document, drawing, projection, three dimensional object, or motion study. Your output can be flat art, contour images and dimensions of solid objects, or data bases for film and video productions, robotics, and more.

SAC sonic digitizers convert graphic information into digital form for processing by host computers or recording and transmission equipment. Processed information can be used to provide cost estimates from drawings, real-time multipoint analysis of motion, robotic instructions, and CAM, CIM, and other automated and semi-automated systems.

SAC systems are used around the world in business, construction, industrial, and medical applications as integrated workstation components or as add-on systems to computer devices.

To Find The Solution To Your Dream.

Call or write today to discuss your special requirements.

SEE US AT COMDEX, BOOTH C746, CAESARS PALACE

feature: the ability to update data in place. Since this allows partial updates of a backup tape, it could potentially save a lot of time. It also lets the system use a tape drive as a block-oriented, random-access medium, which is an especially useful feature in systems where inexpensive but not superfast mass storage is required.

However, proponents of DDS argue that using a tape as a random-access device—starting and stopping it rather than

streaming data smoothly—may promote tape wear. They also claim that writing a driver for such a device is likely to be more difficult than for a streaming device. Is it worth it? The answer is unclear today, but it may make a difference to some applications in the long run.

All the helical-scan drives offer read-after-write data verification and error correction, and they can correct an error detected during a write without breaking stride.

Can QIC Catch Up?

Makers of helical-scan drives claim that they can pack data more densely than on any equivalent "linear" drive because a helical-scan drive doesn't need a multi-track head, or a single one that must shift its position from track to track, to cover almost the entire surface area of the tape. (Doing either, they claim, would introduce mechanical tolerances that limit data density.)

They also claim that the combined motion of the tape and head, which generates an effective head-to-tape velocity of roughly 150 inches per second, provides higher data transfer rates than can be achieved with a linear drive. A linear drive couldn't move the tape at that speed without generating enough friction to melt it!

Meanwhile, QIC adherents are claiming that QIC-1350, a new development standard based on special heads, 900-oersted tape, and data compression, will soon let quarter-inch cartridges store 1.35 gigabytes at data transfer rates of 600K bytes per second. However, the point may soon be moot. Nearly all the major manufacturers of QIC drives have licensed the DDS technology, and many are expected to release DDS drives in the near future.

If they do so, the cost advantages of a mass-market mechanism and the expense of supporting yet another standard may cause them to abandon quarter-inch cartridges for their high-end products. But even if this occurs, quarter-inch cartridges are likely to remain dominant in lower-capacity applications for a long time to come.

The good news for all of us, however, is that mass production and standardization are forcing down prices—at least on the low end of the tape backup spectrum. It's now possible to purchase an internal QIC-40 tape backup unit for around $300 from any of several mail-order dealers. Hopefully, this will be enough to convince most users that there's no sense leaving one's data in danger. ■

L. Brett Glass is a freelance programmer, author, and hardware designer residing in Palo Alto, California. He can be reached on BIX as "glass."

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

Circle 56 on Reader Service Card

M A Y

BIX gives in to special interest groups.

First it was Mac mavens, followed by IBM-ers. Then it was game fans, writers, and our friends, the Amiga aficionados. In less than a year, BIX has established "Exchanges" for each of these large BIX subscriber groups — so people from all over the world can get together to learn, exchange ideas, and simply enjoy themselves.

If you were a BIX subscriber, you could choose from this growing list of Exchange offerings: (In each Exchange, conference name appears on left, description on right.)

Amiga Exchange

amiga.user	Exchange ideas, solve problems, compare notes
amiga.sw	Amiga programming and developer issues
amiga.hw	Amiga hardware design, use, and hookup
amiga.arts	Artistry using the Amiga
amiga.int	Developing for the international Amiga
amiga.special	Special guests and events
amiga	BIX Archive: The original Commodore Amiga conference
amiga.update	Commodore's private conference for certified developers
amiga.com	Commodore's private conference for commercial developers
amiga.dev	Commodore's conference for developers

IBM Exchange

ibm.pc	The venerable PC
ibm.at	The AT series and workalikes
ibm.ps	The PS/2 series
ibm.os2	OS/2 operating system
ibm.dos	PC/DOS & MS/DOS operating systems
ibm.os.386	Alternative 386 operating systems
ibm.utils	Utility software for IBM computers
ibm.repairshop	Garage and Tune-up Shop
ibm.new.prods	New products for IBM computers
ibm.exchange	IBM Exchange clearinghouse
ibm.listings	An index to all Listings files in the IBM Exchange
ibm.other	Applications, printers, modems, etc.
microsoft	Products from Microsoft

Mac Exchange

mac.apple	The word from Cupertino
mac.business	Macs in the office
mac.desktop	Publishing with a Mac
mac.external	Information from all over
mac.hack	Technical information about all aspects of the Mac
mac.macappda	Macintosh Applications Developer's Association
mac.news	Up-to-the-minute information
mac.novice	For beginners
mac.products	Listings of new hardware and software
mac.sandbox	For off-hours fun

Writer's Exchange

desktop.pub	Using microcomputers for publishing
journalism	Reporting and writing news
journalism.pro	Interaction for working press only
marketing	Promotions, sales, public relations, and high tech
new.writers	Getting started in the writing business
poetry.pros	Writing both types of English
sf	For Science Fiction, Star Trek, and fantasy fans
sfwa	Private meeting place for the Science Fiction Writers of America
tech.news	Discuss Microbytes news, product reports, and items
word.processor	Word-processing programs
write.fiction	On writing fiction
writers	The original writers conference
write.fiction	How to write fiction
writers.pros	Interaction for professional writers only
writers.talk	Insights and conversation from professional writers

Interactive Game Exchange

compu.games	About computer games, critiques, tips, and more
dd	Digital Dungeoning: programs and applications and gaming
encounters	The place of all lands and places in space and time
eq	ElfQuest (role-playing games in the land of elves)
ff	The Fantasy Foundation (role-playing games)
fun.n.games	Fun, games, group activities
gazebo	Exploring the on-line social life of BIXen
lexicon	About words
lore	Fantasy-related poetry, prose, song, and art
lworld	Ledinworld (advanced Dungeons & Dragons)
mnm	Meade & Mirth (free-form games set in the Middle Ages — or far in the future)
sca	Society for Creative Anachronisms
snl	Suits & Ladders. By invitation only
snl.preview	Previews what can be found in the closed 'snl' conference

Finally. An on-line service that doesn't nickel and dime you.

It's BIX's flat-fee service

BIX is short for BYTE Information Exchange. The on-line information service that's yours for an unheard-of flat fee of just $39 for three months*—an amount you could easily waste in just two to four hours with an *hourly rate*, on-line service. (Not to mention the fact that you'd be nickel-and-dimed for its monthly minimums.)

And here's another distinction: BIX is strictly for microcomputer pros; it contains no "fluff." As a subscriber, here's what you've got coming to you:

☐ All the information and ideas exchanged in more than 150 microcomputer-related conferences—a give-and-take in which you can participate.

☐ *Microbytes Daily*—up-to-the-minute industry news and new product information.
☐ Plus support from hardware vendors and software publishers, access to extensive software libraries, and the use of our electronic mail service—which allows *binary attachments*.

Subscribe to BIX right now—using your computer and modem

Set your telecommunications program for full duplex, 7 bits, even parity, 1 stop bit. Call BIX directly on our special registration-only number: 800-225-4129 (in Massachusetts, call 617-861-9767). Then hit the return key, and respond as follows:

Prompt:	You Enter:
login (enter"bix"):	bix
Name?	bix.flatfee

You can charge your BIX subscription to major credit cards, or have it billed to your company.

You may also purchase unlimited off-peak access via Tymnet for just $20 per month, or $3 per off-peak hour, in the continental US.** For more information, and your local Tymnet access number, call 800-227-2983 (in NH and outside the US, call 603-924-7681).

*Based on a $156 annual fee, billed quarterly. You may cancel at any time. If you prefer, you may subscribe for a 3-month trial at just $59. Fee does not include Tymnet connect charges.

**No extra charge for 2400-baud access. Tymnet prices are subject to change. International networks access BIX at NUI 310690157800.

BIX

One Phoenix Mill Lane
Peterborough, NH 03458
800-227-2983. In NH 603-924-7681.

IN ANY EVENT

**Pressing a button
on a mouse may seem
trivial, but in
the world of windows,
it's a major event**

Not so long ago, the gauge of a programmer's capabilities was how many boxes of punched cards he or she could carry unassisted. Writing a program was straightforward. Programs possessed a simple, linear overall structure: Your program read its input from the cards, did its calculations, and printed the output.

I suppose the genesis of this was the old batch station—the card reader and line printer combo that you found in the corner of the computer room. Not much

user interaction there beyond pressing the "formfeed" button on the printer when your job was done. The teletypewriter was a rare (and noisy) addition to the computer room and seldom used for more than an alternate portal for input when the batch station was too crowded.

The world has changed. Although I'm sure there are still some card readers around here and there, you don't find many hooked to the latest 386 clones or Mac IIs. Teletypewriters have become little more than sound-effects devices placed offstage during news programs.

The input-process-output structure has also given way to more-complex designs. The avenues by which a program must be able to receive data have multiplied. Now, instead of a card reader, applications use keyboards, mice, serial ports, and—in multitasking systems—messages from other tasks. And nowhere is this diversity of input devices more evident than in the realm of graphical user
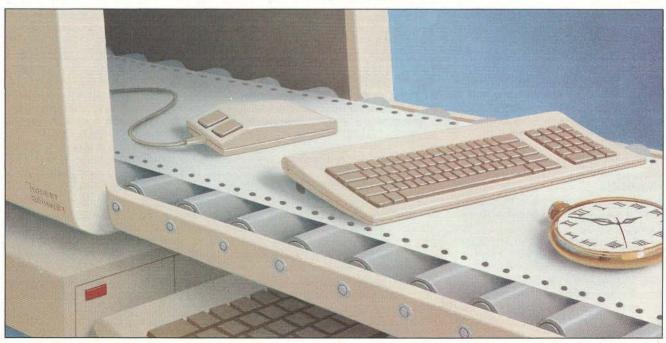
interfaces (GUIs).

In an effort to relieve the snarl of source code that would surely result from a program's having to loop and poll every conceivable point of input, many GUI systems channel input into a single pipeline—the event queue. With this design, the most active code in a program is the looping procedure that retrieves events from the event queue and responds appropriately.

Line 'Em Up

User input occurs asynchronously. That is a fancy way of saying "unpredictably." Of course it's unpredictable—there's a human on the other end of the mouse. Operating-system designers might cope with this situation by providing a library of interrupt service routines—one routine for each device or source of input. Application builders could then hook into the appropriate ISRs and take things as

Listing 1: *C pseudocode for a simple event loop. Messages carrying event information come in through the* GetNextEvent() *routine. The* switch *selects the appropriate handling routine. In this case, the program manipulates a single window.*

```
for(;;)
{
    GetNextEvent(&message);
    switch(message.type)
    {
        case MOUSE:
            DoMouseMsg(message);
            break;
        case KEYBOARD:
            DoKeyboardMsg(message);
            break;

        <...other message types...>

        case QUIT:
            DoExit();
    }
}
```

they came. Coping with interrupts at that level, however, would leave problems such as buffer management in the laps of the application developers.

Instead, the designers of Windows, Presentation Manager (PM), and others have built systems that funnel the asynchronous events into a queue. The system kernel captures the interrupts, packages the information, and places it on a kind of software conveyer belt: the event queue. The application program pulls the event message off the queue, examines it, and then executes the appropriate routine. Voilà: Asynchronous interrupts become an ordered list of events.

(I'm simplifying things. I do that at the beginning a lot. The application program doesn't always just remove events from the queue; it often places them on the queue. Also, more than interrupts triggered by user input will summon an event. Messages coming from other processes can appear as an event as well.)

Programming in the presence of an event queue has encouraged programmers to mold a distinctive process structure around whatever statement fetches the next item from the queue. This structure is called the *event loop*. Although the details vary from language to language (and from system to system), you build an event loop with three primary components: a big "repeat forever" statement, the "fetch next event" function, and some sort of dispatch statement (usually a switch statement). Listing 1 is a pseudocode fragment that shows how the whole thing is built.

The queue is nothing more than a two-ended stack: The system is busy pushing events onto the far end while your program is busy pulling them off the near end. The event queue handling function—next_event()—pulls the next event off the queue and then returns the event's information in a data structure whose contents describe the associated event. For example, if the event is a mouse-click, the data structure would carry pointer coordinate information.

The switch statement simply passes the event data to the appropriate routine. Hence, the do_mouse() routine takes care of mouse functions, the do_keypress() function handles keyboard input, and so on.

As you might guess, things are certainly more complicated than I have shown. A program built around an event queue must keep track of its current state. In a paint program, for example, a mouse-click at coordinates x,y would produce one result if a drawing-tools window is open, and a completely different result if a fonts window is open. Applications typically manage such situations by keeping track of which windows are currently active, where those windows are on the screen, and other such status information.

Multitasking systems allow tasks to operate independently of one another. The event manager maintains a separate event queue for each task. This makes its job more difficult than in a single-tasking system, since now the event manager must know which windows are attached to which task so that events in a given window will go to the proper parent task.

It sounds odd, but there are actually instances when it makes sense to have your program send a message to itself. For example, suppose your paint program has defined a particular key sequence as being a signal to hide its tools window. If the user activates that key sequence, the actual event received will be a key press. So, your program turns around and sends a "hide window" message to itself, with the event message appropriately garnished to indicate that the targeted window is the tools window.

The Object
Another common element of event-driven systems is their handling of complex structures as objects. More specifically, elaborate data structures—such as might define a window—are given handles. Generally, the handles are integers that act as aliases for the structure. (Usually, the handle is a pointer to a pointer, as I described in my August 1989 column. This permits the object referenced

by the handle to be relocatable, a critical asset in graphical systems where windows are always coming and going.)

The use of handles is particularly important with the X Window System, since it is designed to operate across a network. For example, given that a program must manipulate a window—whose defining data structure occupies hundreds of bytes of storage—a single operation could require substantial network traffic if the window's entire data must be exchanged between the client process and the server process. Associating the window with a 32-bit handle means that a request across the network needs only a long integer to identify the window.

In DESQview, the use of objects reduces the complexity of coding the event loop. The object queue structure lets a program treat all sources of input identically and channel their events through a single event queue. The specific event-handling subroutine must resolve the handle into its associated object. The result is a kind of information-hiding: Only the routines that manipulate the object have any idea of the object's internals. DESQview carries this even further; it isn't even necessary for a program to know the details of an object's data structures. A program operates on an object through DESQview calls, never by directly reading or writing data fields within the object's definition.

Macintosh
The Macintosh's event-handling software is divided into two parts. The Mac OS event manager is the lowest (i.e., closest to the hardware). It is here that the Macintosh captures key presses and mouse-clicks. The Mac OS event manager passes events up to the Toolbox event manager, which is what most applications converse with.

Although your program will almost always call the Toolbox event manager, there are cases where you may need to call the Mac OS event manager. Most important, when your program is first launched, it can call the OS event manager's FlushEvents to clear the event queue of any stale events left on the queue from any previous programs. Additionally, the OS event manager's PostEvent is the means by which an application can send events to itself.

An application receives event information in the form of an event record. An event record is composed of

- the event type;
- the time that the event was posted

Up to date.
Down to earth.

changing the world. UNIX is changing the world of computers, the world of business — quite simply, changing the world. s revolutionizing office automation. It's required for U.S. government computer contracts. s the backbone of information strategies worldwide.

the information you need. at's why you need *UNIXWORLD* — e magazine that keeps you to date on the rapidly changing world of open systems mputing. Each issue brings you the latest product trends and technical advances that can fect your business. The inside ory on some of the world's

biggest high-tech companies. Easy-to-understand programming tips and tutorials that can help you and your company use UNIX to its fullest. And unbiased hardware and software reviews to help you invest wisely when you buy.

The whole UNIX-verse.
UNIXWORLD'S in-depth features go beyond dry technical facts, to show how the pieces fit together — to tell you what's important about the advances and the strategies that are changing your world. And *UNIXWORLD* consistently offers the freshest, most down-to-earth writing you'll find in any computer publication.

into the event queue;
- the location of the mouse when the event took place;
- the state of the mouse button and the modifier keys (the Shift, Option, and clover keys) when the event occurred; and
- information that is dependent on the type of the event.

The last entry is referred to in Macintosh documentation as the *message field*. It's a long integer whose contents differ from event type to event type. For example, if the event is a key press, the lowest byte of the message field contains the ASCII code corresponding to the key. The next-higher byte (bits 8–15) carries the scan code of the key pressed. As you continue reading, you'll discover that most of the other systems I'll mention use an event record (sometimes called a message) structure that's a variation on the Macintosh theme.

The Mac provides a filter for incoming events called the *event mask*. Each bit of the event mask corresponds to a particular type of event. Setting a bit in the event mask lets the corresponding event type

"pass through" to the application; clearing a bit blocks the associated event type. Therefore, if your application is—heaven forfend!—command-line driven and thus has no need of the mouse, you can filter out mouse events so that your program needn't bother with them.

The Mac's core event-handling routine is GetNextEvent. It takes two arguments: the event mask word and a pointer to a memory buffer where the next event record will be stored. GetNextEvent returns a Boolean value indicating whether there's an event waiting for your application to handle. It also performs some background work that your application doesn't explicitly see. For example, Get-NextEvent checks whether an alarm that you've set on the alarm clock should be triggered; it also activates any routines that have been attached to command-shift-number keys (FKEY functions).

The Mac's event queue can handle up to 20 events. If the queue fills up, the system pushes old events into oblivion to make room for incoming new events. The GetNextEvent routine actually pulls events off the queue. Sometimes, however, your program may wish to only

examine events on the queue without removing them. In this case, you call the EventAvail routine, which operates just like GetNextEvent without extracting events from the queue.

Keep in mind that the event manager is an application's sole source of user input. The Mac provides a journaling feature: Items passing through the event queue can be fed to an audit file while on their way to the application. Then, at some later time, you can put the program into "playback mode," in which the contents of the audit file are fed directly to the event queue. This lets you create quite complex demonstration programs that include mouse movements and button clicking, as well as keyboard input.

Microsoft Windows
The Microsoft Windows world is a little different from that of the Mac—an understatement if ever there was one. But if you look beyond the obvious differences in host processors and parent companies, Windows distinguishes itself by virtue of its multitasking capabilities.

Windows actually maintains two event
continued

---

queues. Your program is really never aware of one, the system queue. As its name implies, the system queue is maintained by Windows and is the system-wide clearinghouse for all events. Your application program actually accesses the other event queue. Since each application gets its own queue, it's up to Windows to determine the recipient of a given event message. Now that I think about it, I suppose that means that a Windows session actually supports multiple event queues; the number is dependent on how many tasks are simultaneously active.

Windows event messages are compact (as on the Mac), and they are all built using the same four fields:

- a 16-bit handle to a window,
- a 16-bit message type,
- a 16-bit parameter, and
- a 32-bit parameter.

The last two fields are the Windows counterparts to the Mac's message field. (Don't get confused by my usage of the word *message*. Windows, and everyone else following, refers to event records as messages.) They carry data specific to the event. For example, if the event indicated a movement of the mouse (defined by Windows as a WM_MOUSEMOVE event message), the 32-bit parameter carries the concatenated x and y coordinates of the mouse.

Windows has its own meaning for *synchronous* and *asynchronous*: A synchronous event message pops out of the queue in the same order (relative to everyone else in the queue) that it went in. However, as soon as an asynchronous event enters the queue, it moves immediately to the front. Agreed, an asynchronous event sounds dangerous, but its raison d'être is to alert a task to some situation that must be handled immediately.
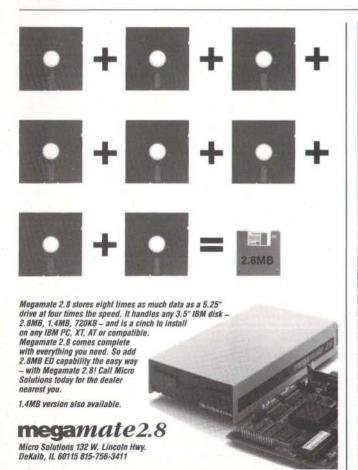
A Windows event loop revolves around three functions: GetMessage() receives an event message from the Windows kernel. The message is passed to TranslateMessage(), where any key-down messages are translated to their appropriate ASCII equivalents. Finally, the message is handed to DispatchMessage(), which delivers it to the proper window.

Of course, I'm oversimplifying. How does a window accept delivery of a message? The answer lies in what occurs in the initialization code that every Windows program executes just after start-up. Specifically, the program must register with the operating system the *kinds* of windows that it intends to use. (This is referred to in Windows and OS/2 as registering the window *class*.)

During the registration process, the program provides each window class with a pointer to a routine that will handle all messages delivered to windows of that class. So when I say that the window accepts delivery, that translates to "the window's associated handling routine is called, and detailed information about the message is passed in as arguments." The switch-statement portion of the event loop is within each window-handling routine.

Listing 2 outlines the registration process. In its initialization portion, the program hands Windows a template of the classes Window1 and Window2. When the program creates instances of those windows (with the CreateWindow() function), Windows knows which templates to use to draw the windows on-screen, as CreateWindow() specifies the class.

continued

Listing 2: *Event handling in a multiwindow environment—in this case, Windows. Notice that the* switch *statement has been decoupled from the* while *statement. The* DispatchMessage() *routine sends the message to the proper window-handling routine—*Wind1Proc() *or* Wind2Proc()*—where message filtering takes place.*

```
/*
** Initialization routine to register
** a window with the operating system.
*/
BOOL FAR PASCAL Initialize(myhandle)
HANDLE myhandle;          /* Handle to myself */
{
WNDCLASS mywind1;
WNDCLASS mywind2;
/*
** Name mywind1's class and associate a
** window-handling routine.
*/
mywind1.lpszClassName="Window1";
mywind1.lpfnWndProc=Wind1Proc;
/*
** Do same with mywind2.
*/
mywind2.lpszClassName="Window2";
mywind2.lpfnWndProc=Wind2Proc;
/* Register with Windows. */
RegisterClass(&mywind1);
RegisterClass(&mywind2);
return;
}
/*
** Create two windows of class mywind1
** and mywind2.
*/
BOOL FAR PASCAL Createwinds(myhandle);
HANDLE myhandle;          /* Handle to myself */
{
HWND hWind1, hWind2;     /* Handles to windows */
/* Create two windows */
hWind1=CreateWindow((LPSTR)"Window1",(LPSTR)"MyW1",...
            myhandle,(LPSTR)0);

hWindw=CreateWindow((LPSTR)"Window2",(LPSTR)"MyW2",...
            myhandle,(LPSTR)0);

<..code to display the windows goes here...>
}
```

```
long FAR PASCAL Wind1Proc(whandle,msgtype,wp,lp)
HWND whandle;             /* Window handle */
unsigned msgtype;         /* Message type */
WORD wp;                  /* Word parameter */
LONG lp;                  /* Long parameter */
{
switch(msgtype)
{
        case WM_ONETHING:
                <code to handle one thing>
                break;
        case WM_ANOTHER:
                <code to handle another>
        default:
                return(DEfWindowProc(whandle,msgtype,wp,lp);
                break;
}
return(0L);
}

long FAR PASCAL Wind2Proc(whandle,msgtype,wp,lp)
HWND whandle;             /* Window handle */
unsigned msgtype;         /* Message type */
WORD wp;                  /* Word parameter */
LONG lp;                  /* Long parameter */
{
< code virtually identical to Wind1Proc above >
}
int Pascal WinMain(myhandle,...)
HANDLE myhandle;
{
MSG mymsg;

        Initialize(myhandle);
        CreateWinds(myhandle);
        while(GetMessage(&mymsg,NULL,0,0))
        {       TranslateMessage(&mymsg);
                DispatchMessage(&mymsg);
        }

        <...termination code goes here...>
}
```

Listing 2 solves another riddle. How does Windows know which program owns a particular window? When you launch an application under Windows, the system assigns the application a unique identifier—myhandle in listing 1. The CreateWindow() function takes myhandle as an argument when a particular instance of a window is created. The argument is the key to how the operating system knows which application is a given window's keeper.

Presentation Manager

The user interface supported by OS/2's PM is the same as that employed by Windows; in fact, PM's application programming interface is based on the Windows API. Consequently, it's no surprise that the two systems share many traits.

As in Windows, applications running under PM respond to event messages passed in from PM itself. Also, PM carries on with the use of handles—16- or 32-bit aliases to complex structures such as windows, message queues, and so on.

Unlike Windows programs, PM applications must explicitly create their own event message queue. (My guess is that this is because OS/2 might be running PM and non-PM applications simultaneously. A non-PM application may not need an event message queue, so programs create their own queues only when necessary.) A program does this with a call of the form

```
msgqhand=
    WinCreateMsgQueue(hab,0)
```

where msgqhand is a handle to the created queue. The variable hab is a handle to what is referred to as an *anchor block* and is used more or less as a program's means for uniquely identifying itself in the PM environment.

PM provides the program with this unique identifier via the WinInitialize() routine, which you must call prior to calling WinCreateMsgQueue(). This anchor block handle is the counterpart to the handle passed into a Windows application when it is launched.

The format of a PM event message looks like a Windows event message with some finery added:

- the handle of the window to which the message is directed;
- a 16-bit message type field;
- two 32-bit parameter fields;
- the time the message entered the queue; and
- the position of the mouse when the message entered the queue.

As you can see, a PM message has a special field for the mouse coordinates. In Windows, whenever the application receives a mouse message, the rodent's coordinates are kept in the 32-bit parameter (see above).

(Aside: PM further deviates from Windows in the nomenclature it applies to event messages. In PM, synchronous messages are referred to as *queued*, asynchronous messages as *nonqueued*.)
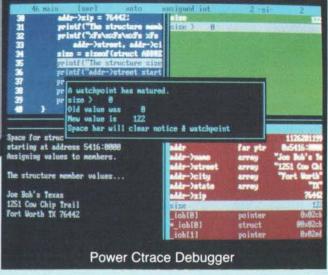
A PM event-loop kernel takes up about two lines of source code. It looks something like this:

```
while(WinGetMsg(...))
    WinDispatchMsg(...);
```

I've removed the functions' arguments. The WinGetMsg() function retrieves a message from the message queue, while WinDispatchMsg() sends that message off to the appropriate destination. Since WinGetMsg() is the argument to the while loop, apparently this loop exits only when WinGetMsg() returns a value of 0. That's exactly what happens when your program receives a WM_QUIT message (sent whenever the user selects Close from the menu).

As in Windows, the switch-statement portion of the event-handling loop is carried in a window's window-handling procedure. Oddly, the time of posting and the mouse-coordinates portion of the event structure are not passed to the window-handling routine—only the handle, message type, and parameter fields are. I

suppose this is done to keep memory use trim; only messages generated by the mouse demand that the application know where the pointer was at the time of the event. Your program can retrieve the mouse location with a call to the Win-QueryMsgPos() function.

X Window System

The X Window System (referred to as X Window hereafter) is a network-based windowing system released by MIT (with significant contributions by DEC and other companies). Since it works across a network, its structure is client-server. User programs run as clients, passing requests to the server, which does most of the grunt work of managing a bit-mapped display. User programs talk to the server through a set of C routines gathered into the Xlib library.

Since X Window can operate over a network, the transformation of interrupts (from key presses, mouse-clicks, and so on) into event messages is mandatory. The hardware that generates user-input interrupts may be several meters' worth of coaxial cable away from the hardware executing the application program.

X Window carries on the theme of using handles to manipulate complex entities. However, a handle is referred to as a *resource identifier*, since abstract objects are called *resources*.

The fulcrum routine of the X Window event loop is XNextEvent, which takes two arguments. Its first argument is a pointer to a display structure, which carries information about the current screen and the X Window server. (This display structure is X Window's way of uniquely identifying a given user. It is the counterpart of OS/2's anchor block and the my-handle variable in Windows.)

The second argument is a pointer to an XEvent structure. This is a union of structures, each member structure of which defines a type of event. The first integer in the XEvent acts as a selector and lets you pick which member of the union to use for a particular event. (If you find this confusing, consider XEvent a data structure whose architecture differs from event type to event type.)

The XNextEvent routine doesn't do any event filtering; you set that with the XSelectInput routine prior to calling XNextEvent. As with the other event-driven systems I've mentioned, the filtering is handled by a mask. In this case, the mask is a 32-bit word.

DESQview

DESQview is another multitasking operating system. Most of the programs that

you execute under it are "DESQview un-aware"; they're simply programs written for MS-DOS that don't know any better. But you can create "DESQview aware" programs that use the multitasking capabilities. (See my February column.)

In DESQview, items that a program communicates with (screens, keyboards, and so on) appear in the form of objects. So your program manipulates window objects (display output), panel objects (dialog boxes, menus, and help screens), keyboard objects (self-explanatory), pointer objects (mice), mailbox objects (semaphores and interprocess communication), and objectq objects (the event queue). I pronounce "objectq" as "object queue" to help remind me of its use. As in the other systems I've mentioned, objects are referenced by handles (32-bit identifiers in DESQview).

The objectq is the foundation for the event loop in a DESQview application. You might reasonably imagine that your program would accept keyboard input by creating a keyboard object and attaching that keyboard object to the current task. Whenever the task expects input, it sends a READ message to the keyboard object. This scenario would apply as well to a task requiring mouse input: The task creates and then opens a pointer object.
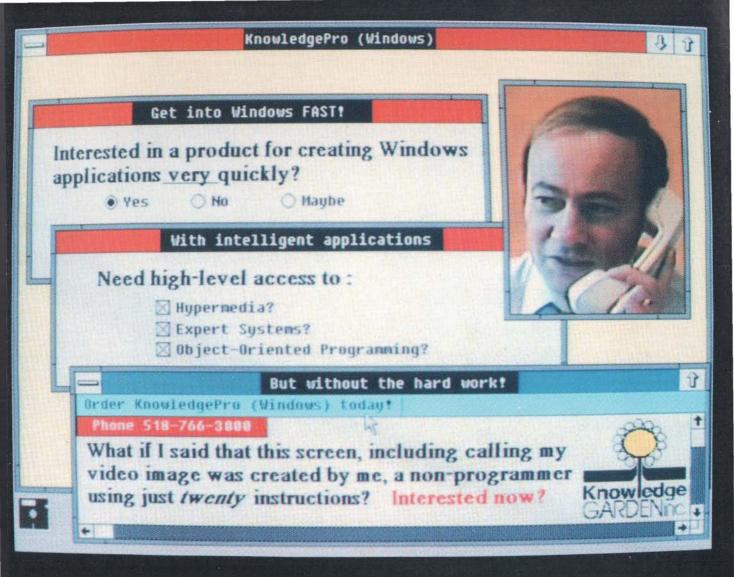
That's not precisely true, and here's where the objectq comes in: Each item on the objectq is a handle to an object for which an event is pending. So, if your program reads the objectq and receives the handle to the keyboard object, your program knows that the keyboard holds the next pending event. This saves your program from having to issue passels of READ messages helter-skelter to all the objects it is manipulating. A DESQview event loop ends up looking something like what I've shown in listing 3.

In addition, the objectq sidesteps the potential for a task's falling asleep at the wheel. Specifically, if a task issues a READ call to a keyboard object when no keyboard input is available, that task is suspended until a key is pressed. So if the user is using the mouse while the task is waiting on the keyboard, mouse input may get lost. (Or worse, the system can simply freeze up.)

For Example

To illustrate event-loop programming, I put together a small calculator program for DESQview. It's simply a hexadecimal-to-decimal conversion calculator. But it gives me the opportunity to showcase some of the special features DESQview provides for user input.

The calculator's front panel consists of 19 keys: the digits 0 through 9, the characters *A* through *F*, and the special keys *, 10, and 16. You can guess the use of many of the number and letter keys. The 10 key sets the display to base 10, 16 sets the display to base 16, and * clears the display. Keys can be activated by pointing and clicking with the mouse, as well as typing their counterparts on the keyboard. (You access 10 and 16 through the F1 and F2 function keys.)

This calculator program takes input from the mouse as well as from the keyboard. You'd therefore expect that the event loop would require some means of distinguishing between messages from the keyboard and messages from the mouse. This is not so, thanks to a component of DESQview referred to as the *field manager*. The field manager lets you make a window "aware" of those parts of itself that make up user input and output.

You activate the field manager by building a data structure called a *field table* and transmitting that table to the window as part of what is referred to as a *window stream*. This is data that you send to a window to alter the behavior of the window's contents (instead of data for display). In this case, the field table's elements describe all the input and output fields within that window—where those fields are and how they act. So the input fields are the calculator's keys; the output field is the display.

Of course, the calculator's keys are really more or less buttons. DESQview calls fields of this sort *select* fields; you

can select them by pressing an associated key from the keyboard or by moving the mouse pointer to the field and clicking the button. Fields that can actually accept arbitrary strings of input are simply called *input* fields.

Once you've defined the floor plan of your window (i.e., where all the fields are located and what their attributes are), assembled that into a field table, and sent that to the window, your work is largely done. Your program can capture all user input by opening a keyboard object attached to the window. The field manager consolidates mouse and keyboard input and sends it to the keyboard object.

Furthermore, information that your program receives from the keyboard object has had considerable intelligence injected into it by the field manager. For example, since the field manager knows where fields are within the window, it attaches the field number to any received input and sends that to your program via the keyboard object. Your application doesn't have to determine which field has been input based on mouse coordinates.

This is true for input fields as well as select fields. In my calculator, I need only know which key was selected. In a complex data-entry screen, I could have fields for clients' names, dates, and addresses. The user could select a field and enter data, and the field manager will report to my program what the user typed, as well as what field he or she was in when it was typed.

Detailed information—mouse buttons and such—can be had from a variable called the *keyboard status*. DESQview

returns this variable whenever you read the keyboard object. This is how the user closes the calculator: A specific bit in the keyboard status field indicates when the user has pressed the Escape key. My program reads an Escape key as a signal to shut down.

Final Event
It's relieving to see that, with all the diversity in computers today, there's still plenty of similarity. I suppose Mac programmers have a leg up on the rest of us; they were dealing with event loops on microcomputers before there was any mention of PM or X Window.

Perhaps, if you look at it right, event-driven systems have been around a lot longer than you may have guessed. It's just that their early forms were subtle. Why, I clearly remember feeling like an object in a queue, waiting for my chance at the card reader and line printer. And when the program worked...what an event *that* was. ∎

Editor's note: The full source code is available in a variety of formats. See page 5 for details.

Rick Grehan is the director of the BYTE Lab. He has a B.S. in physics and applied mathematics and an M.S. in computer science/mathematics from Memphis State University. He can be reached on BIX as "rick_g."

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

BYTE

PRODUCT SHOWCASE

- ■ **BUYER'S MART**
- ■ **BYTE BITS**
- ■ **PRODUCT SPOTS**
- ■ **MICRO PRODUCT CENTER**

THE BUYER'S MART
A Directory of Products and Services

THE BUYER'S MART is a monthly advertising section which enables readers to easily locate suppliers by product category. As a unique feature, each BUYER'S MART ad includes a Reader Service number to assist interested readers in requesting information from participating advertisers.

Effective January 1, 1990.
RATES: 1x—$590 3x—$550 6x—$525 12x—$475 24x—$450
Prepayment must accompany each insertion. VISA/MC Accepted.

AD FORMAT: Each ad will be designed and typeset by BYTE. Advertisers must furnish typewritten copy. Ads can include headline (23 characters maximum), descriptive text (250 characters is recommended, but up to 350 characters can be accommodated), plus company name, address and telephone number. Do not send logos or camera-ready artwork.

DEADLINE: Ad copy is due approximately 2 months prior to issue date. For example: November issue closes on September 8. Send your copy and payment to THE BUYER'S MART, BYTE Magazine, 1 Phoenix Mill Lane, Peterborough, NH 03458. For more information call Brian Higgins at 603-924-3754.

ACCESSORIES

SIMMS DIRECT FROM MFG.
*HP LASERJET SERIES II MODULE**PS/2 70 & 80 MODULES*
*LASERWRITER NTH MODULES**MACINTOSH MODULES*
*IBM COMPATIBLE MODULES**AST MODULES*
2 YR WARRANTY
SAVE 60% BUYING DIRECT AND RECEIVE A *FREE* FLOPPY DISK HOLDER AS WELL. MFG. IS ALSO AVAILABLE USING YOUR MEMORY CHIPS. CALL NOW FOR MEMORY NEEDS.
BEPHYR INDUSTRIES, INC.
PH: (714) 951-5193   FAX: (714) 951-1541

Inquiry 576.

NEW! WET RIBBON RE-INKER
Stop discarding expensive printer ribbons. Our Wet Ribbon Re-inker gives you more than 50 better-than-new ribbons for each one you buy!
- Easy to use
- Accommodates most cartridges
- High quality synchronous motor
- Automatic timer
- Satisfaction guaranteed!
BLUE RIBBON INK, LTD.
6803 E. 47th Ave. Dr., Unit D, Denver, CO 80216
(800) 477-3465   **(303) 333-8880**

Inquiry 577.

CUT RIBBON COSTS!
Re-ink your printer ribbons quickly and easily. Do all cartridge ribbons with just one inker! For crisp, black professional print since 1982. You can choose from 3 models:
Manual E-Zee Inker — **$39.50**
Electric E-Zee Inker — **$94.50**
Ink Master (Electric) — **$189.00**
1000s of satisfied users. Money-back guarantee.
BORG INDUSTRIES
525 MAIN ST., JANESVILLE, IA 50647
1-800-553-2404   Fax: 319-987-2251

Inquiry 578.

80C31 DEVELOPMENT BOARD
Space for 32K Eprom   2K Scram   8 Bit Latched Outputs
12 I/O Lines   RS232 Serial Port   5V Only—All CMOS Components
Can be used for development of microcontroller-based projects. Currently used as the base controller for: display driver, remote scoreboard, serial/parallel converter, thermal printer controller, transmission line buffer, cartridge storage device and many more!
$99 WITHOUT EPROM — $109 WITH EPROM
CUSTOM FIRMWARE AVAILABLE
COMPUTER ADVICE
515 Calle San Pablo, Camarillo, CA 93010
(805) 482-2002

Inquiry 579.

COMPANION AND EXTENDER
Place a keyboard and monitor up to 600' from your CPU with EXTENDER and COMPANION products. Keep a second Keyboard/Monitor at the CPU with COMPANION. Supports MDA, CGA, EGA, VGA, PS2. Uses single ⅜" cable.
Prices start at $149.00 for EXTENDER and $219.00 for COMPANION 25 ft. unit complete.
CYBEX CORPORATION
2800-H Bob Wallace, Huntsville, AL 35805
205-534-0011   International Fax #205-534-0010

Inquiry 580.

ACCESSORIES

FREE CATALOG
A complete source for all your computer supplies — media, paper, cables, furniture, software, ribbons, laser, cleaning & FAX supplies, accessories & much more.
Bulk Diskettes — Minimum Order 50

5.25" DS/DD	25¢	3.5" DS/DD	65¢
5.25" DS/HD	49¢	3.5" DS/HD	$1.19

GAAN COMPUTER SUPPLIES
186 B East Sunnyoaks Ave., Campbell, CA 95008
(800) 523-1238, In Calif. (408) 370-6747

Inquiry 581.

Finally a Better Toner Cartridge for your Canon® PC Copier: HP® or Apple® Laser Printer
REFILL KIT IS $19.95   SAVE NEARLY $250.00 on average
or 3 FOR $50.00   three refills over 4 new cartridges
New Modified* Toner Cartridges to accept refill kits (*holes are pre-drilled)

PC 3-5 Cartridge	$89.95	HP & Apple Series I, EP, CX $109.95
PC 6,7 Cartridge	$129.95	HP & Apple Series II, EPS, SX $109.95
PC 10-25 Cartridge	$99.95	many other Laser Printers that use Canon Engines

These toner cartridges are modified to easily be refilled up to *three additional times* with our inexpensive do-it-yourself kit. Or use your own empty cartridge and modify yourself with simple instructions in refill kit. Colors also available.
Call: **Morack, Inc.**
9132 Windsor Dr., Palos Hills, IL 60465
Phone: (708) 598-0580   1-800-837-9696   Fax: (708) 598-9203

Inquiry 582.

Printer Buffer: $99.95
256 Kbyte buffer controls complete printing function. Never wait for your printer again. Simply activate buffer and return to any computing task while files are printing.
Serial, parallel, and converter models available to 2 Megabytes (512 Kbyte: $119.95; 1 Megabyte: $159.95). Liberal 30 day return policy assures satisfaction!
(513) 644-2230
Technologic Systems
18277 Timber Trails • Marysville, Ohio 43040

Inquiry 583.

HEWLETT PACKARD
Buy — Sell — Trade

Laser Jet II/IID	1400/2400	Color Pro (7440)	550
Genuine HP 2 Meg/4 Meg	400/1200	HP-7550A	1700
Desk Jet	400	Draft Pro DXL/EXL	3300/5000
Rugged Writer	800	Draftmaster I/II	5500/6000
Electrostatic Plotters	C1600 (D Size)/C 1601 (E Size)		14000/16500

Science Accessories Corporation Sonic Digitizers
36" x 48" (2750)   60" x 72" (3175)
T. E. Dasher & Associates
4117 Second Ave. S., Birmingham, AL 35222
Phone: (205) 591-4747 Fax: (205) 591-1108

Inquiry 584.

APPLICATION GENERATOR

VERY EASY C PROGRAMMING
Includes complete Unix, Microsoft and Turbo C compatible source code.
1. CG: Creates C program from data   ...$179
2. CGW: Creates Windows programs   ...$189
3. CGX: Creates C++ programs   ...$169
4. HD: Programmable Windows Hypertext   ...$159
Other products include Input Processor $129, Prolog interpreter $79. Call for complete catalog. No source versions and demo disk are also available.
A. I. Coder
32651 N. Burr Oak, Solon, OH 44139
(216) 349-4850

Inquiry 585.

ARTIFICIAL INTELLIGENCE

MANA
Intelligent Computing for 286/386
Transform your PC into an AI-based workstation. Expert systems, AI, knowledge base applications, rapid prototyping, object oriented, multi-tasking, virtual memory, LISP/PROLOG etc. modes in multi-language compiler, mouse driven, multi-windowed, fully integrated environment.
BMS, Inc.
700 E. Ogden Avenue, Westmont, IL 60559
Tel: (708) 789-9160   Fax: (708) 789-0253

Inquiry 586.

NATURAL LANGUAGE C LIBRARY
Increase your market share! Use JAKE to add a natural language front end to your application. JAKE translates English queries and commands into C function calls and data structures. JAKE offers context-sensitive semantic processing; interfaces easily; <64K mem.
JAKE $495. INTERACTIVE DEMO $10
ENGLISH KNOWLEDGE SYSTEMS, INC.
5525 Scotts Valley Dr. #22, Scotts Valley, CA 95066
(408) 438-6922

Inquiry 587.

NanoLISP $99.99
An MS-DOS Common LISP interpreter that supports most Common LISP operations and strictly adheres to the standard. Numerous advanced and extra features, excellent debugging facilities, sample AI programs, fully-indexed manual, free technical support.
Microcomputer Systems Consultants
P.O. Box 6646, Santa Barbara, CA 93160
(805) 967-2270

BAR CODE

LABELING SOFTWARE
On EPSON, IBM, OKI dot matrix or LaserJet. Flexible design on one easy screen. Any format/size. Up to 120 fields/label. 18 text sizes to 3"-readable at 100'. AIAG, MIL-STD, 2 of 5, 128, UPC/EAN, Code 39. File Input & Scanned logos/symbols (PCX)—$279. Other programs from $49. 30-day $$ back.
Worthington Data Solutions
417-A Ingalls St., Santa Cruz, CA 95060
(800) 345-4220   In CA: (408) 458-9938

BAR CODE READERS
For PC, XT, AT, & PS/2, Macintosh, and any RS-232 terminal. Acts like 2nd keyboard, bar codes read as keyed data. With steel wand—$399. Top rated in independent reviews. Works with DOS, Xenix, Novell, Alloy, -ALL software. Lasers, magstripe, & slot badge readers. 30-day $$ back.
Worthington Data Solutions
417-A Ingalls St., Santa Cruz, CA 95060
(800) 345-4220   In CA: (408) 458-9938

THE BUYER'S MART

AT BUS DESIGN

At last, here is the **timing book** for the XT and AT Bus. Detailed text, tables and diagrams tell you what each signal line is for, what it does and when it does it. All the information is compatible with the IEEE P996 Specification for the ISA (AT) Bus. In addition, the 8 and 16 bit parts of the **EISA** Bus are included. **AT Bus Design,** by Ed Solari, has over 200 pages, with more than 100 figures and tables. Handy 7" x 9" format, soft cover, $69.95.

FREE

We'll include a **free** copy of the pocket-sized **XT-AT Handbook by Choisser and Foster** with each AT Bus Design book if you tell us where you saw this ad. Of course, this $9.95 value is also available by itself. Or buy five or more for only $5.00 each.

800-462-1042
In California 619-271-9526

MasterCard **VISA**

Annabooks
12145 Alta Carmel Ct., Suite 250
San Diego, CA 92128 **Money-back guarantee**

AccSys™ for C PARADOX

- Set of C libraries for easy access to Paradox data and index files.
- No need to concern yourself with internal files format.
- Increased performance shortens processing time.
- Total control over table files primary and secondary index files.
- Network and single-user versions available.

Copia International Ltd.

1964 Richton Drive, Wheaton, Illinois 60187
708/682-8898 FAX: 708/665-9841
Instant info. via FaxFacts: 617/740-0025 Press: 1625007#

VOICE MASTER KEY® SYSTEM II

VOICE RECOGNITION & SPEECH RESPONSE

FOR IBM PC/XT/AT/386, PS/2, LAPTOPS, COMPATIBLES

FOR PRODUCTIVITY, PRESENTATIONS, SOFTWARE DESIGN, ENTERTAINMENT, LANGUAGE TRAINING, EDUCATION, MORE...

SPEECH/SOUND RECORDING AND PLAYBACK. Desktop Audio sound editing allows you to create custom sound applications. Variable sample rate (to 20 KHz) and compression levels. A four-voice music synthesizer is included also!

VOICE RECOGNITION TSR utility allows you to add voice command keyboard macros to your CAD, desktop publishing, word processing, spread sheet, or entertainment programs. Up to 64 voice commands in RAM at once—more from disk.

HARDWARE SYSTEM contains built-in speaker with separate volume and tone controls, external speaker and headphone jacks. Enclosure made of sturdy vinyl-clad steel. Attaches to parallel printer port without affecting normal printer operation (U.S. Patent 4,812,847). Headset microphone, printer cable, 9 volt AC adapter (110 volt UL/CSA listed), and comprehensive user manual included.

QUALITY THROUGHOUT. MADE IN USA. ONLY $219.95

ORDER HOTLINE: (503) 342-1271 Mon-Fri, 8 AM to 5 PM PST

Visa/MasterCard, company checks, money orders, CODs (with prior approval) accepted. Personal checks subject to 3 week shipping delay. Specify computer type when ordering. Add $5 shipping charge for delivery in USA and Canada. Foreign inquiries contact Covox for C&F/CIF quotes. OEM configurations available.

30 DAY MONEY BACK GUARANTEE IF NOT COMPLETELY SATISFIED.

CALL OR WRITE FOR FREE PRODUCT CATALOG

COVOX INC.
675 Conger Street
Eugene, Oregon 97402

TEL (503) 342-1271
FAX (503) 342-1283
BBS (503) 342-4135

DISC DRIVE REPAIR COUPON

Formatted Cap.	Flat Rate	SPECIAL
10-19 mb	$99	89.10
20-29 mb	$125	112.50
30-39 mb	$150	135.00
40-49 mb	$175	157.50
50-85 mb	$210	189.00
86-120 mb	$275	247.50
121-150 mb	$325	325.00
151-275 mb	$425	425.00
276-380 mb	$495	495.00
TEST & EVALUATION $25		

SHIPPING YOUR DRIVE FOR REPAIR
Pack your drive carefully and well protected in a sturdy shipping box. Include with the shipment a note with your name, address, daytime telephone number and a brief description of the problem with the drive. If prepaying, allow $9 for shipping and insurance costs.

DATA RECOVERY - CALL FOR QUOTE

FLOPPY REPAIRS
5.25" & 3.5" • $45   8" • $135

VALID THROUGH 6/30/90
COUPON MUST BE SENT WITH DRIVE

DISC DRIVE SALES

SPECIAL OF THE MONTH

42.8 mb, 28/35 ms Avg. Access Time **$295**

Auto Parking, 6.5 W Power, 3.5" w/5.25", Frame Available
80g Non Operating Shock Rating, 1 Year Warranty

MACINTOSH SCSI SUBSYSTEMS	HARD CARDS
HD/SCSI - 20.....................$295	10MB/85MS Hard Card $185
HD/SCSI - 30.....................395	20mb/65ms Hard Card 225
HD/SCSI - 40.....................495	30mb/65ms Hard Card 275
HD/SCSI - 80.....................895	40mb/65ms Hard Card 310
HD/SCSI - 150....................1195	48mb/36ms Hard Card 395

THOUSANDS OF DISC DRIVES IN STOCK — We Feature Technical Support for Everything We Sell — We Specialize in Disc Drives — Ask for Our Brochure

jb TECHNOLOGIES, INC.
5105 Maureen Lane Moorpark, CA 93021

Phone: 805/529-0908
Fax: 805/529-7712

Circle 155 on Reader Service Card

BACK ISSUES FOR SALE

BYTE
Limited Quantities — Order Now!

	1987	1988	1989	1990
January	■			
February		■		
March				
April				
May	■			
June				■
July				■
August				■
September				■
October				■
November				■
December	■			■
Inside the IBM PCs	■			■

Available Issues ☐

Rates (postage and handling included):

1987–'90 BYTE Issues	$6.00*
BYTE '83–'84 Index	$4.00
BYTE 1985 Index	$4.00
BYTE 1988 Index	$4.00
1985 Inside the IBM PCs	$4.00
1986 Inside the IBM PCs	$5.00
1988 Inside the IBM PCs	$6.00
1989 Inside the IBM PCs	$6.00

*June 1988 (Benchmarks) **$3.00**
*December 1988 **$3.00**

The above prices include postage in the U.S. Please add $.50 per copy for Canada and Mexico; and $2.00 per copy to foreign countries (surface delivery). European customers please refer to Back Issue order form in International Advertising section of book.

Please indicate which issues you would like by checking (√) the boxes.

Send requests with payments to:

BYTE Back Issues
One Phoenix Mill Lane
Peterborough, NH 03458
(603) 924-9281

☐ Check enclosed. Payments from foreign countries must be made in U.S. funds payable at a U.S. bank.

Name _____

Address _____

City _____ State _____ Zip _____

Charge: ☐ VISA ☐ MasterCard

Card # _____

Exp. Date _____

Signature _____

All orders must be prepaid. Please allow four weeks for delivery.

Buy with
Confidence

In an effort to make your telephone purchasing a more successful and pleasurable activity, The Microcomputer Marketing Council of the Direct Marketing Association, Inc. offers this advice, "A knowledgeable buyer will be a successful buyer." These are specific facts you should know about the prospective seller before placing an order:

Ask These Important Questions

- *How long has the company been in business?*
- *Does the company offer technical assistance?*
- *Is there a service facility?*
- *Are manufacturer's warranties handled through the company?*
- *Does the seller have formal return and refund policies?*
- *Is there an additional charge for use of credit cards?*
- *Are credit card charges held until time of shipment?*
- *What are shipping costs for items ordered?*

Reputable computer dealers will answer all these questions to your satisfaction. Don't settle for less when buying your computer hardware, software, peripherals and supplies.
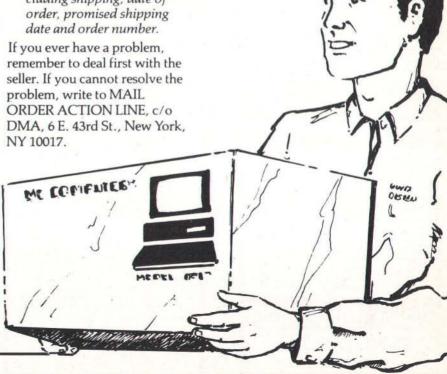
Purchasing Guidelines

- *State as completely and accurately as you can what merchandise you want including brand name, model number, catalog number.*
- *Establish that the item is in stock and confirm shipping date.*
- *Confirm that the price is as advertised.*
- *Obtain an order number and identification of the sales representative.*
- *Make a record of your order, noting exact price including shipping, date of order, promised shipping date and order number.*

If you ever have a problem, remember to deal first with the seller. If you cannot resolve the problem, write to MAIL ORDER ACTION LINE, c/o DMA, 6 E. 43rd St., New York, NY 10017.

This message is brought to you by:

the MICROCOMPUTER MARKETING COUNCIL of the Direct Marketing Association, Inc.
6 E. 43rd St.,
New York, NY 10017

MMC
MICROCOMPUTER
MARKETING COUNCIL
of the Direct Marketing Association, Inc.

EDITORIAL INDEX BY COMPANY

Index of companies covered in articles, columns, or news stories in this issue
Each reference is to the first page of the article or section in which the company name appears

Company, Page #   Inquiry #

READER SERVICE

To get further information on the products advertised in BYTE, fill out the reader service card by circling the numbers on the card that correspond to the inquiry number listed with the advertiser. This index is provided as an additional service by the publisher, who assumes no liability for errors or omissions.

* Correspond directly with company.

Alphabetical Index to Advertisers

READER SERVICE

*Correspond directly with company.

READER SERVICE

To get further information on the products advertised in BYTE, fill out the reader service card by circling the numbers on the card that correspond to the inquiry number listed with the advertiser. This index is provided as an additional service by the publisher, who assumes no liability for errors or omissions.

* Correspond directly with company.

Index to Advertisers by Product Category

READER SERVICE

*Correspond directly with company.

REQUEST <u>FREE</u> INFORMATION BY FAX

Attention BYTE Readers!! Now you can fax your requests for free product and advertiser information featured in this issue.

Just fax this page to 1-413-637-4343. You'll save time because your request for information will be processed as soon as your fax is received.

1 Circle the numbers below which correspond to the numbers assigned to advertisers and products that interest you.

2 Check off the answers to questions "A" through "C".

3 Print your name, address, and fax number clearly on the form.

4 Remove this page or copy this page clearly and fax it to the number above.

Fill out this coupon carefully. PLEASE PRINT.

Name

Title

Company

Address

City

State/Province   Zip

Country

(   )   (   )
Phone Number   Fax Number

A. What is your level of management responsibility?
1 ☐ Senior-level Management
2 ☐ Other Management
3 ☐ Non-Management

B. What is your primary job function/principal area of responsibility? (Check one.)
4 ☐ Administration
5 ☐ Accounting/Finance
6 ☐ MIS/DP/Information Center
7 ☐ Product Design and Development
8 ☐ Research and Development
9 ☐ Manufacturing
10 ☐ Sales/Marketing
11 ☐ Purchasing
12 ☐ Personnel
13 ☐ Education/Training
14 ☐ Other:

C. Please indicate your organization's primary business activity: (Check one.)

Computer-Related Businesses:
15 ☐ Manufacturer (Hardware, Software)
16 ☐ Computer Retail Stores
17 ☐ Consultants
18 ☐ Service Bureau/Planning
19 ☐ Distributor/Wholesaler
20 ☐ Systems House/Integrator/VAR
21 ☐ Other:

Non-Computer-Related Businesses:
22 ☐ Manufacturing
23 ☐ Finance, Insurance, Real Estate
24 ☐ Retail/Wholesale
25 ☐ Education
26 ☐ Government
27 ☐ Military
28 ☐ Professions (Law, Medicine, Engineering, Architecture)
29 ☐ Consulting
30 ☐ Other Business Services
31 ☐ Transportation, Communications, Utilities
32 ☐ Other:

☐ *I subscribe to BYTE.*   ☐ *I do not subscribe to BYTE.*
☐ *Please send me one year of BYTE Magazine for $24.95 and bill me. Offer valid in U.S. and possessions only.*

BYTE McGraw-Hill   MAY 1990

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40
41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60
61	62	63	64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79	80
81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	100
101	102	103	104	105	106	107	108	109	110	111	112	113	114	115	116	117	118	119	120
121	122	123	124	125	126	127	128	129	130	131	132	133	134	135	136	137	138	139	140
141	142	143	144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159	160
161	162	163	164	165	166	167	168	169	170	171	172	173	174	175	176	177	178	179	180
181	182	183	184	185	186	187	188	189	190	191	192	193	194	195	196	197	198	199	200
201	202	203	204	205	206	207	208	209	210	211	212	213	214	215	216	217	218	219	220
221	222	223	224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239	240
241	242	243	244	245	246	247	248	249	250	251	252	253	254	255	256	257	258	259	260
261	262	263	264	265	266	267	268	269	270	271	272	273	274	275	276	277	278	279	280
281	282	283	284	285	286	287	288	289	290	291	292	293	294	295	296	297	298	299	300
301	302	303	304	305	306	307	308	309	310	311	312	313	314	315	316	317	318	319	320
321	322	323	324	325	326	327	328	329	330	331	332	333	334	335	336	337	338	339	340
341	342	343	344	345	346	347	348	349	350	351	352	353	354	355	356	357	358	359	360
361	362	363	364	365	366	367	368	369	370	371	372	373	374	375	376	377	378	379	380
381	382	383	384	385	386	387	388	389	390	391	392	393	394	395	396	397	398	399	400
401	402	403	404	405	406	407	408	409	410	411	412	413	414	415	416	417	418	419	420
421	422	423	424	425	426	427	428	429	430	431	432	433	434	435	436	437	438	439	440
441	442	443	444	445	446	447	448	449	450	451	452	453	454	455	456	457	458	459	460
461	462	463	464	465	466	467	468	469	470	471	472	473	474	475	476	477	478	479	480
481	482	483	484	485	486	487	488	489	490	491	492	493	494	495	496	497	498	499	500
501	502	503	504	505	506	507	508	509	510	511	512	513	514	515	516	517	518	519	520
521	522	523	524	525	526	527	528	529	530	531	532	533	534	535	536	537	538	539	540
541	542	543	544	545	546	547	548	549	550	551	552	553	554	555	556	557	558	559	560
561	562	563	564	565	566	567	568	569	570	571	572	573	574	575	576	577	578	579	580
581	582	583	584	585	586	587	588	589	590	591	592	593	594	595	596	597	598	599	600
601	602	603	604	605	606	607	608	609	610	611	612	613	614	615	616	617	618	619	620
621	622	623	624	625	626	627	628	629	630	631	632	633	634	635	636	637	638	639	640
641	642	643	644	645	646	647	648	649	650	651	652	653	654	655	656	657	658	659	660
661	662	663	664	665	666	667	668	669	670	671	672	673	674	675	676	677	678	679	680
681	682	683	684	685	686	687	688	689	690	691	692	693	694	695	696	697	698	699	700
701	702	703	704	705	706	707	708	709	710	711	712	713	714	715	716	717	718	719	720
721	722	723	724	725	726	727	728	729	730	731	732	733	734	735	736	737	738	739	740
741	742	743	744	745	746	747	748	749	750	751	752	753	754	755	756	757	758	759	760
761	762	763	764	765	766	767	768	769	770	771	772	773	774	775	776	777	778	779	780
781	782	783	784	785	786	787	788	789	790	791	792	793	794	795	796	797	798	799	800
801	802	803	804	805	806	807	808	809	810	811	812	813	814	815	816	817	818	819	820
821	822	823	824	825	826	827	828	829	830	831	832	833	834	835	836	837	838	839	840
841	842	843	844	845	846	847	848	849	850	851	852	853	854	855	856	857	858	859	860
861	862	863	864	865	866	867	868	869	870	871	872	873	874	875	876	877	878	879	880
881	882	883	884	885	886	887	888	889	890	891	892	893	894	895	896	897	898	899	900
901	902	903	904	905	906	907	908	909	910	911	912	913	914	915	916	917	918	919	920
921	922	923	924	925	926	927	928	929	930	931	932	933	934	935	936	937	938	939	940
941	942	943	944	945	946	947	948	949	950	951	952	953	954	955	956	957	958	959	960
961	962	963	964	965	966	967	968	969	970	971	972	973	974	975	976	977	978	979	980
981	982	983	984	985	986	987	988	989	990	991	992	993	994	995	996	997	998	999	1000
1001	1002	1003	1004	1005	1006	1007	1008	1009	1010	1011	1012	1013	1014	1015	1016	1017	1018	1019	1020
1021	1022	1023	1024	1025	1026	1027	1028	1029	1030	1031	1032	1033	1034	1035	1036	1037	1038	1039	1040
1041	1042	1043	1044	1045	1046	1047	1048	1049	1050	1051	1052	1053	1054	1055	1056	1057	1058	1059	1060
1061	1062	1063	1064	1065	1066	1067	1068	1069	1070	1071	1072	1073	1074	1075	1076	1077	1078	1079	1080
1081	1082	1083	1084	1085	1086	1087	1088	1089	1090	1091	1092	1093	1094	1095	1096	1097	1098	1099	1100
1101	1102	1103	1104	1105	1106	1107	1108	1109	1110	1111	1112	1113	1114	1115	1116	1117	1118	1119	1120
1121	1122	1123	1124	1125	1126	1127	1128	1129	1130	1131	1132	1133	1134	1135	1136	1137	1138	1139	1140
1141	1142	1143	1144	1145	1146	1147	1148	1149	1150	1151	1152	1153	1154	1155	1156	1157	1158	1159	1160
1161	1162	1163	1164	1165	1166	1167	1168	1169	1170	1171	1172	1173	1174	1175	1176	1177	1178	1179	1180
1181	1182	1183	1184	1185	1186	1187	1188	1189	1190	1191	1192	1193	1194	1195	1196	1197	1198	1199	1200
1201	1202	1203	1204	1205	1206	1207	1208	1209	1210	1211	1212	1213	1214	1215	1216	1217	1218	1219	1220
1221	1222	1223	1224	1225	1226	1227	1228	1229	1230										

CHAOS MANOR MAIL

Jerry Pournelle answers questions about his column and related computer topics

High-Priced Floppy Disks

Dear Jerry,

As a user of 3½-inch high-density disks, I am frustrated by the monumental prices charged for them. I fail to understand how these disks can cost two to three times as much as the 3½-inch standard-density disks when the only apparent difference is the magnetic media.

To make matters even worse, recent ads for a glorified hole punch—which converts standard disks to high-density disks—seem to indicate that there is no difference in the two media at all. What gives? Are we being bilked?

I used to shake my finger at people who lost data because they didn't back up their data frequently. At $100 per backup, all I can say is, "I know how you feel."

Roy Anderson

I get conflicting stories about those disks, but I have concluded the following:

At least with all the major media makers (e.g., Sony and Dysan), all the disk media is made by the same process and in the same factories. The product is then tested. Some disks fail when tested as high-density, although most don't. The ones that don't fail are marked high-density and put in containers with the high-density hole punched in them. This process continues until there are sufficient high-density disks to fill the marketing demand for them.

The rest of the disks are then subjected to less stringent tests—as double-sided double-density disks, for example—and those that pass are packaged accordingly until the market demand for them is met. This goes on until all the demands, including those for single-sided single-density disks, have been met.

Note what has happened: It may be that no disk ever failed any test. On the other hand, some have not been tested at any data density higher than that marked on the package. Moreover, depending on yields and demands, some disks may in fact have failed a higher-density test, and some disks sold as single-sided may have failed a double-sided test.

I have no data on what defects can cause a disk to fail in one mode but not in another (other than the obvious: one side may be flawed). If there are statistics on this, the disk companies are not publishing them.

Conclusion: You take a chance when you buy one of those hole punches and make a high-density disk out of a standard double-sided double-density disk. I cannot calculate how dangerous that is. I don't do it, and I don't recommend that you do. The best I can accomplish in buying cheap disks is to save a couple of hundred dollars. The worst that can happen is to lose data worth—in my case, at least—a lot more than that.

A last warning: Some disks have not been very highly burnished, and thus they have rougher media than others. It's certainly true that such disks can cause extra head wear or deposit bits of the media on the disk head, thus requiring more frequent head cleaning. For all those reasons, I buy quality premium disks; the brand I prefer is Dysan, largely because I've been using their disks since 1976 and I've had no reason to change. I don't doubt that there are other premium disks as good.—Jerry

Nota Bene, Pro and Con

Dear Jerry,

In your reply to Joseph M. Prospero's letter (February 1989), you mentioned having recommended Nota Bene to "a young person writing a dissertation."

I believe I am that person. I'm writing to let you know that the dissertation is finished and that Nota Bene worked like a dream. In fact, I wouldn't have finished without it. The program requires too much memory (which is in short supply because I wrote chapters that were as long as 275K bytes), but it's worth every byte. I used Nota Bene's Special Language Supplement to enter diacritical marks for romanized Japanese and Sanskrit (e.g., Ōkubo Dōshū, Ākāsagarbha, and maṇḍala). Can you imagine trying to write these in by hand (accurately!) in a 500-page manuscript?

The most wonderful thing about Nota Bene, however, is its speed and ease of use. It really is a dream marriage of simplicity and versatility. This is one program whose advanced features get used because they are easy to use.

William Bodiford
Tarzana, CA

I think 275K-byte chapters are self-defeating, but then I'm not on your degree committee. I am glad that everything worked out well for you.—Jerry

Dear Jerry,

I read with interest the letter from Joseph M. Prospero, recommending that you move up to Nota Bene instead of considering XyWrite III Plus. I think Nota Bene may well be the software marketing bomb of the decade.

As Prospero says, Nota Bene is a customized version of XyWrite. Unfortunately, it has been "customized" in such a way that, from the user's point of view, it functions as a totally different word processor. I work in a university, and my job is to edit academic texts. For this we use XyWrite, which is compatible with the Atex system used to publish the texts. Some academic departments bought Nota Bene because they need its additional features, such as Hebrew and Greek. That's how I got to try Nota Bene and compare it to XyWrite. I thought it would be easy—because the programs are fundamentally the same, it should be possible to use both, XyWrite for the straightforward jobs and Nota Bene for the specialized ones. But to use Nota Bene would require almost as much retraining as changing to Microsoft Word, WordPerfect, or one of the other word processing programs.

Both XyWrite and Nota Bene have far more logical command structures than WordPerfect, but XyWrite uses the Alt key for word functions, while Nota Bene uses the Control key for word functions and the Alt key for paragraph functions. In XyWrite, F9 is the Execute key, while in Nota Bene it is F10, and so it goes on.

Why did the companies do this? The
continued

only reason I can think of is that either XyQuest insisted on it, or Dragonfly got the idea of doing it so it wouldn't be competing for the same territory. Whatever the reason, it is a great pity that they did it. XyQuest and Dragonfly had a winning combination, and with it they could have swept the academic and publishing worlds at least. Instead, they've painted themselves into a corner. People will continue to buy WordPerfect, with its difficult user interface, because one difficult interface is probably easier to learn than two easy but incompatible ones.

Steve Hayes
Pretoria, South Africa

Well, people tend either to love or hate Nota Bene, which, now that you remind me, is the case with XyWrite, too. —Jerry

Update Syndrome

Dear Jerry,

I just finished reading your column in the July 1989 BYTE, and it led me to the following thoughts about the relationship between software updates and those people who are afflicted by them.

There are very few updates that I look forward to. The only one that comes immediately to my mind is my annual notice from SoftView that the new version of MacInTax is ready. Of course, SoftView knows that updates are at the heart of its business, so they do them right.

Some update notices appear determined to make me feel technically archaic. I've received update notices for MacDraw and MacProject telling me that I need a Mac Plus to run them. I called Claris and told them that I have a 512KE with a memory upgrade. "Sorry," they said, "you need to have the Plus." Well, the demo packages that my clients let me try work fine on my old system. Maybe Claris just wants to hold down the number of upgrades.

Upgrades can provide wonderful hardware restrictions to accompany wonderful software features. When I received an upgrade to FileMaker, it refused to run on my Mac. I called the company, and the person I spoke with told me that the problem was that I didn't have a SCSI port on my system. "But your announcement said that it would run on a floppy disk drive–only system," I said. A few days later, a special version

of the program arrived that didn't check for the SCSI.

Update announcements don't tell you what you need to know. Last week I got a notice that I could upgrade to Borland's Reflex 2.0 at a very attractive price. As a registered user of Reflex for the Mac, I thought it was strange that the order card offered the choice of a 3½-inch or 5¼-inch disk. There was nothing in six pages of update advertising that described the hardware needed to run it.

I can't figure out who these updates are for. In only one case (a little utility called KiwiEnvelopes) has a company written to me asking about how I use its software and what sorts of things I might like to see in an improved version. Do other companies try to find out what their users are actually doing, or do they just check the features in their competitor's press releases?

I'm being held hostage by my software. I've just moved up to version 4 of Microsoft Word for the Mac. It doesn't give me any major benefit over version 3.02 for the things I do, but I can't shake the feeling that if I call up next January with a problem I'm having with 3.02, I'll find out that support is available for version 4 only.

Most updates don't make the good stuff better. Where are the improvements to the things that made me buy the program in the first place? There's always plenty of "new" stuff, and I'm sure that there are people to whom it's very important. I'm not one of them. The new "table" feature in Word looks neat, but an improved spelling checker would be more useful.

On another subject, I've looked at the announcements and reviews for the Sinclair, and I'm still having difficulty seeing the degree to which it's an improvement over my old Tandy 102, particularly with the Traveling Software ROM. Does it all come down to the quieter keyboard?

John Boddie
Newark, DE

I agree completely with your update blues. As for the Sinclair, it runs a long time on not much power, and it's a lot lighter than the Tandy. And quieter keys really do make a difference. —Jerry ∎

Jerry Pournelle holds a doctorate in psychology and is a science fiction writer who also earns a comfortable living writing about computers present and future. He can be reached c/o BYTE, One Phoenix Mill Lane, Peterborough, NH 03458, or on BIX as "jerryp."

PRINT QUEUE

Hugh Kenner

Does the Walking Do the Fingers?

*How the great mathematicians did it,
and how the telephone does it to us*

Over and over, in the 53 years since Lancelot Hogben's book *Mathematics for the Million* made a stir by selling and selling for months on end, the expositor's formula of choice has been historical. Start way back when it all began, with the Egyptians and their knotted ropes. Then tell lots of stories. That's all fine until the stories start obstructing the foreground.

In his *Men of Mathematics* (1937), E. T. Bell solicits our tears over Évariste Galois, age 20, scribbling by candlelight all night before a fatal duel, but he never does try to make us see why the scribblings matter. (Worse, it now turns out that the famous last night is almost all fiction. Galois didn't invent group theory against the clock. He did write a few letters. The story had gotten romanticized before Bell; then Bell's inability to read the French sources abetted his ideologized desire to create a tormented radical snuffed out young.)
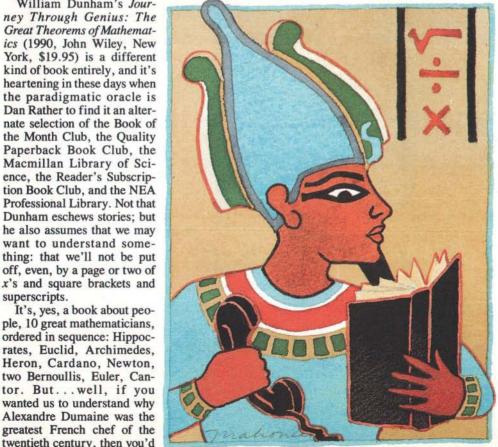
William Dunham's *Journey Through Genius: The Great Theorems of Mathematics* (1990, John Wiley, New York, $19.95) is a different kind of book entirely, and it's heartening in these days when the paradigmatic oracle is Dan Rather to find it an alternate selection of the Book of the Month Club, the Quality Paperback Book Club, the Macmillan Library of Science, the Reader's Subscription Book Club, and the NEA Professional Library. Not that Dunham eschews stories; but he also assumes that we may want to understand something: that we'll not be put off, even, by a page or two of x's and square brackets and superscripts.

It's, yes, a book about people, 10 great mathematicians, ordered in sequence: Hippocrates, Euclid, Archimedes, Heron, Cardano, Newton, two Bernoullis, Euler, Cantor. But...well, if you wanted us to understand why Alexandre Dumaine was the greatest French chef of the twentieth century, then you'd serve us a meal of Dumaine's, which is now impossible; whereas Dunham can serve us a theorem of Heron's (aka "Hero of Alexandria"), who achieved the general formula for a triangle's area although fettered by impossibly primitive notation. Heron called on five prior propositions, three of them Euclid's, and achieved "one of the best surprise endings in mathematics." For, "as with a good Agatha Christie novel, we can be within a few lines of the end and still have no idea how the matter will be resolved. Yet we need not fear, for he ultimately brings the strands together in a wonderful climax."

A man who sets an exposition up like that may possibly be setting himself up for a pratfall. Dunham, though, doesn't falter. Again and again, on that pattern, he leads us through the preliminaries to a proof, through the proof itself, and then—rightly—invites us to marvel. Always, his attention is on the reasoning power someone brought to bear with the then-available tools.

With the then-available tools. That means, there's an unwavering historical focus. So we're shown Euclid working toward a proof of the Pythagorean theorem—which is the climax of book I of the *Elements*—via necessarily sequenced propositions. The result—square on the hypotenuse, you know it—"was well known before Euclid's day." But Euclid wanted a proof he could make with "a rather lean tool-kit"—his postulates, "common notions," and a mere 46 already-proven propositions.
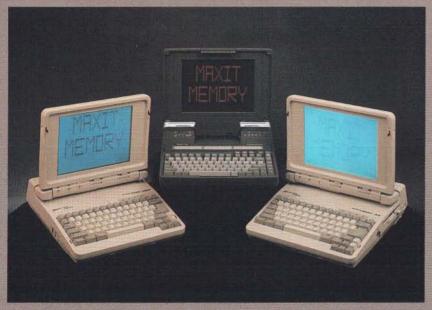
"It is surely possible to devise short proofs of the Pythagorean theorem by using similar triangles," but those won't come up till book VI, and Euclid clearly wanted to get to it "as early and directly as possible." Also, having no algebra—$a^2 + b^2 = c^2$—he had to work with plain drawings. His great proof, out of context, has been made to seem clumsy. But by the time Dunham has put in place all

those easily forgotten restrictions, the only response that fits his climactic exposition is awe.

He's equally good on Newton, at age 23, inventing the binomial theorem, in free time he had while Cambridge was shut down by plague; then five years later using it to evaluate pi. From a mere nine-term series, he got seven correct decimal digits; later, from 20 terms, he got 16. ("I am ashamed to tell you," he added, "to how many places of figures I carried these computations, having no other business at the time.")

What Avital Ronell would have us understand is the world a century with the telephone has lodged us in.

The anecdotal details have been reported often. What I can't display in this space is something else, the way William Dunham illuminates the inner workings of the pi theorem itself. That's a rare compliment he's paying his lay readers.

Lay readers will (I hope) also open Avital Ronell's *The Telephone Book*, if only to go "Arrrrgh!" before sticking around to be hooked. (*The Telephone Book: Technology, Schizophrenia, Electric Speech*, 1990, University of Nebraska Press, Lincoln, $35.) An ongoing, if tacit, theme of *Journey Through Genius* is the difference a mental environment can make; base-10 numbers with zero for placeholder created a mind-set unimaginable to Euclid and Archimedes. Euler, writing x's and y's, extracting roots, may never have reflected on the extent to which he and his greatest predecessors lived in utterly different worlds.

What Avital Ronell would have us understand is the world a century with the telephone has lodged us in. The blurb says she worked with "an advanced form of optical fiber," and it takes a moment to realize the meaning of that: She just *looked*, mostly at texts.

From the primal soup on her first pages, we can pick out one-liners like "The telephone . . . presupposes the existence of another telephone, somewhere" and "America operates according to the logic of interruption and emergency calling." Also, when you answer the phone, you accept the role of "automatic answering machine." (Dammit, the phone just rang, and my hand reached. Is Ms. Ronell on to something? What? Well, nothing ever prodded Newton's hand to just reach, his mind to regress into answer mode.)

A son of mine helps edit a magazine in New York. As a frequent contributor, I've frequent need to call him. With every call, I'm aware of interrupting something. I've grown aware, too, how that's become the rhythm of his psyche: mental work between random interruptions. His next caller may be a lawyer in Newark, a flack in Soho. In the interstices, he crafts his resourceful prose. The phone is simply part of his environment, like his desk lamp. I myself, having never reached a truce with it, am happiest when my wife screens the calls.

For let me declare two biases. Having lost much of my hearing at age 5, I'm more estranged from the telephone than are most of you. Also, for a number of years, from age 26, I was intimate with Marshall McLuhan, later (after I'd lost touch with him) the media guru. So I came at *The Telephone Book* from a special angle.

It has black covers, it's twice as tall (10 inches) as it's wide, and it sports real Yellow Pages, devoted to Notes and Index. Here and there, type on a line gets bigger and smaller, even gets illegible (simulating line noise?), and folk that Marshall would have called "linear thinkers" will run away screaming. Moreover, its target reader will be up on Heidegger, Derrida, Freud, deconstruction, and arguments that pivot on puns; also deft at fielding German and French citations.

Just paging through the first half tried my patience. Then at page 227 my attention locked. For Mr. Thomas J. Watson—no, not IBM's fabled Old Man, but the "Come here! I want you!" Watson of the Bell story—left us a nearly unknown autobiography, where we learn of a queer time (1875) when the world contained but two telephones and no telephone system.

The Thing had begun life as a pure idea: "speech-shaped electric current," something any competent electrician (unlike A. G. Bell) knew was impossible. And once they'd made it, what *was* it? So little did it meet a felt need, it stirred up fear, contempt, blushes. "When a prominent lawyer tried it," and "heard my voice on the telephone making some simple remark, he could only answer after a long embarrassed pause, 'Rig a jig, and away we go.'"

Bell and Watson thought of it as an obstreperous child with a will of its own. Disadvantaged Thing, "it talked better nights and Sundays," when "the diminution of city sounds gave the telephone a much better chance to be heard." Notice that trope: Like E.T., the wee newcomer is trying to make itself heard. Tom Watson was a dabbler in the occult, and getting the child to talk sense could resemble a séance. (One time, without Bell's knowledge, he even consulted a medium. No dice.)

Word of it attracted people who "heard voices"—schizophrenics. One visitor knew how to talk at *any* distance without a wire. He did it all the time; even offered (for 15 dollars a week) to let Watson take off the top of his skull and check circuitry.

And the first long-distance call, Boston to East Cambridge, amounted to "Ahoy! Ahoy!" "Are you there?" "Do you hear me?" "What's the matter?" Watson found random noise more interesting: "I used to spend hours at night in the laboratory listening to the many strange noises in the telephone and speculating as to their cause." There were snaps and chirps that he attributed to "explosions on the sun" or "signals from another planet." He may have been, as Ms. Ronell remarks, "the first convinced person ever actually *to listen to noise*."

Gone days. Now we casually pick up The Thing, or move binary data with its aid, or dial to order matinee tickets on a charge card. What I'm recommending of *The Telephone Book* is just (in samplings) the second half. I know I'll be going through that second half again, to seek more nuances in The Thing's initial strangeness. For ponder, "Are you there?" Where's there? There's nobody there. That voice isn't in the room, it's in your head: a statement that, prior to 1875, was never applicable to a voice, unless you were a schizophrenic. ∎

Hugh Kenner is a professor of English at Johns Hopkins University. His reviews have appeared in publications like the New York Times *and* Harper's. *His recent books include* A Sinking Island *and* Mazes. *He can be contacted on BIX as "hkenner."*

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

SETTING THE STANDARDS

15 YEARS

RETURN OF THE COLOSSAL CODE

Sherman, set the wayback machine for 1975

B ack in my old data-processing days—just as the flame was being held to the fuse of the microcomputer explosion—I worked for a Memphis company that specialized in accounting and record-keeping packages for radio and TV stations. I was tucked into a two-man office with a fellow named Chuck, who had already been working there for some months. The nameplate on Chuck's desk didn't carry his name; instead, it read: "I'm so confused." At first, I thought this was just a joke.

Then I began trying to do some programming for the company. These people used walls of Data General minicomputers and big, washing-machine-size hard disk drives. I discovered that all work was done in assembly language buttressed by gigantic macro routines, each of which carried on for several lines—sometimes pages. Deciphering each routine's function was often impossible. Worse, as I dug deeper into what source code I could find, the number of routines proliferated exponentially. It took me days just to get a program to say "Hello, world" on my terminal.

Chuck came to my rescue. He had gone through the same learning shock that I was in the middle of, and he had put together a catalog of all the macro routines and what they did. "Don't worry,"

he said, "I've already made a dictionary of all that. Here." My desk thundered as he dropped the foot-thick homemade manual in front of me. Don't worry? Geez, I thought, is there a concordance for this thing? I'll be buried for weeks just looking for the table of contents!

Years later, I'm in my office at BYTE when one of the editors walks in carrying his proud new bundle: a back-cracking box of disks and manuals. My desk thunders as he drops it in front of me. It is a C compiler. Not a whole operating system or an illustrated encyclopedia, just a C compiler. Remember C—the language that Kernighan and Ritchie described in that thin white book with the pale blue capital *C* on the front cover? This monster on my desk is so big that the software company can't ship it on 360K-byte floppy disks anymore; they'd need about 30 disks if they tried. They use 1.2-megabyte floppy disks.

The sight of such a huge box for a C compiler makes me dizzy. I still have my old copy of the first CP/M manual that I ever purchased. Some issues of BYTE are thicker. Yet I got plenty done with CP/M: I wrote programs, edited files, deleted files, and did some word processing. I could cope with it.

Two shelves up from that old CP/M manual is a whole shelf of red AT&T Unix manuals. I have two Unix machines that I use frequently: I write programs and edit and delete files, and I even once tried to do some word processing. The kinds of things that I do in Unix are not substantially different from the kinds of things I did in CP/M. Yet there's all this extra *stuff* I have to figure out—starting friction, I suppose—just to get some work done.

Before it begins to look like I've singled out Unix, I'll mention the time when a representative from Microsoft demonstrated the benefits of OS/2 Presentation Manager's on-line manual. It was a ride through screenfuls of functions, functions, and more functions, and it left me

feeling like I'd just had my turn on Space Mountain. What amazed me the most wasn't just the amount of source code you needed to print "Hello, world" in a window; it was everything you had to know *before* you could even begin to write a "Hello, world" program.
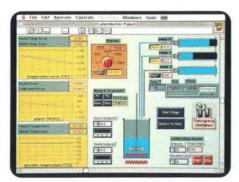
Don't get me wrong; I understand that complex jobs require complex tools. A CP/M machine certainly wouldn't do well as a network server. But I wonder sometimes if we're not injecting complexity where it isn't needed. Manipulating the computer becomes more involved than the actual work we need to do to complete whatever job we're tackling. Remember those early programs of the cassette-driven Apples, VICs, and TRS-80s? Remember how you were going to throw out your index cards and keep your family's dinner recipes on the computer?

I see similar things happening today. I see people using pop-up TSR programs for phone lists and appointment logs—spending hours tweaking this and that to get the format just right for the morning printout—when a Rolodex and a desk calendar would probably do just as well. I once spent three days with a project system, printing out Gantt charts and flow diagrams for every staff member in the BYTE Lab for the next two months. A week later I sat down with some notebook paper and a pen and did much the same thing in about 10 minutes. (I'll admit that I messed up a few times and had to cross stuff out, but I'll redraw everything when I get a free moment.)

I have this nightmare. We're all up on the crest, riding the wave of new microcomputer technology. But then the wave washes us up on a beach. And there stand all the minicomputer and mainframe guys with thick COBOL manuals and piles of source code listings. They look down at us and say: "Where've you been? We've been waiting for you." ■

Rick Grehan is director of the BYTE Lab. He can be reached on BIX as "rick_g."